

HP StorageWorks Continuous Access EVA administrator guide

Part number: T3687-96019
2nd edition: May 2005



Legal and notice information

Copyright © 2003-2005 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

UNIX® is a registered trademark of The Open Group.

Java™ and Solaris™ are trademarks of Sun Microsystems, Inc.

Red Hat® and Red Hat® Enterprise Linux are registered trademarks of Red Hat, Inc.

Linux® is a registered trademark of Linus Torvalds.

Printed in the U.S.A.

HP StorageWorks Continuous Access EVA administrator guide

Contents

Preface	9
About this guide	9
Intended audience	9
Prerequisites	9
Related documentation	9
Document conventions and symbols	11
HP technical support	12
HP-authorized reseller	12
Helpful web sites	12
Providing Feedback	13
1 About HP Continuous Access EVA	15
Overview	15
Overview of HP Continuous Access EVA	15
Features	16
Setup and configuration assumptions	17
Licensing	17
Hardware and software components	17
Array	17
Hosts and host operating systems	18
Virtual Controller Software	18
Management server and software	19
Interface options	20
HP Command View EVA Tasks	22
HP Replication Solutions Manager Tasks	23
2 Concepts	25
Overview	25
Local replication	25
Snapclone	25
Snapshot	25
Remote data replication	25
Bidirectional replication	26
DR groups	26
Source and destination DR groups	26
Replication direction	26
Home designation	27
DR group presentation	27
Presenting DR group members to the same HBA	28
DR group properties	28
DR group log	28
DR group log states	28
DR group log size	29
Managed sets	29
Members	29
Managed set properties	30
Managed set actions	30
Considerations	30
General considerations	30
Storage resource considerations	30

DR group considerations	30
Failover	30
Considerations	31
Replication manager logs	33
Failsafe mode	33
3 Recovery	35
Overview	63
Planning for a disaster	35
Failover	36
Backing up configuration information	37
Using SSSU to capture configuration information	37
Manually capturing configuration information	38
Possible event scenarios	39
Performing recovery actions	40
Planned failover	40
Unplanned failover	44
Resume operations if unable to access destination while source in failsafe-locked state (extended period of time)	41
Return operations to Home array	41
Return operations to replaced new storage hardware	41
Disk group hardware failure on the source array	41
Disk group hardware failure on the destination array	41
Failover and recovery procedures	42
Planned failover	42
Unplanned failover	44
Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)	44
Revert to Home (failback)	45
Return operations to replaced new storage hardware	46
Recovering from a disk group hardware failure	49
Failed disk group hardware indicators	49
Disk group hardware failure on the source array	49
Disk group hardware failure on the destination array	52
4 Operating System Specifics	61
Overview	63
Stopping host I/O	55
Allowing remote host discovery of devices	55
HP OpenVMS	55
HP Tru64 Unix	56
HP-UX	56
IBM AIX	57
Linux	57
Microsoft Windows NT	57
Microsoft Windows 2000/2003	57
Novell NetWare	58
Sun Solaris	58
5 Troubleshooting	61
Overview	63
LUN inaccessible to host	61
DR groups in unknown state	61
Tunnel thrash	61
Remote server cannot detect a destination LUN	62
Long delays or time-outs on HP-UX	62

6 Best practices	63
Overview	63
Creating a destination snapclone before making a full copy	63
Data movement using a snapclone	64
Manually specifying disk group membership for a log-EVA 3000/5000 only	65
Three-site cascaded data replication using snapclones	66
Before you begin	68
Procedure	69
Post procedure	69
Bootless DR group planned failover with Linux using LVM in standalone mode or with	
SuSE SLES 8 running LifeKeeper 4.4.3	70
Source host procedure	70
Destination host procedure	71
Red Hat and SuSE Linux Lifekeeper clusters	71
Throttling of merge I/O after logging	72
Backing up replication jobs and configurations	72
Optimizing discovery refresh intervals	72
Optimizing browser-based GUI performance	72
Coordinating enabled-host downtime	72
Minimizing simultaneous jobs	72
Avoiding configuration changes while jobs are running	72
Optimizing the number of active enabled hosts	73
Coordinating enabled host shutdowns	73
Coordinating replication server shutdowns	73
Avoiding network identification changes	73
Maintaining network connections	73
Using log files for troubleshooting jobs	73
Making CD-ROMs of replication product Web download files	73
Managing replication events	74
Minimizing simultaneous replication events on an array	74
Avoiding simultaneous replication events for the same virtual disk	74
Job scheduling	74
Complying with EVA snapshot rules	74
Complying with EVA snapclone rules	74
Caching in Microsoft Windows	75
 Glossary	 77
 Index	 81

Figures

1 Basic HP Continuous Access EVA configuration with redundant servers	15
2 HP EVA management environment	18
3 HP Continuous Access EVA DR group replication	27
4 Replicating relationships among DR groups	32
5 Planned and unplanned transfer of operations	43
6 Resumption of operation if unable to access destination in failsafe mode	45
7 Array names having failed or new hardware (destination) and your surviving array (source)	47
8 Disk Group Hardware Failure window	50
9 Data Replication Group Manual Deletion window	51
10 Vdisk Deletion in Progress window	51
11 Disk Group Hardware Failure window	53
12 Data Replication Group Manual Deletion window	53
13 Vdisk Deletion in Progress window	54
14 Creating a DR group from a snapclone	65
15 Data movement using snapclones example	67
16 Add a Host window	68
17 Operation succeeded page	69

Tables

1 Document conventions	11
2 Other interfaces and products that perform replication	20
3 When to and when not to fail over a DR group, managed set, or array	37
4 Manual configuration form	38
5 Array log	46
6 Replication manager display icons	49

Preface

About this guide

This guide provides the following information about:

- HP StorageWorks Continuous Access EVA product features
- Network configuration and replication concepts
- Event monitoring
- Failover and recovery procedures
- Best practices and troubleshooting

Intended audience

This guide is intended for system and network administrators who are experienced with the following:

- Storage area network (SAN) configurations
- Host operating system (OS) environments
- Enterprise Virtual Arrays (EVAs), referred to as arrays in this guide

Prerequisites

Before using this document, read and follow the information in [Setup and configuration assumptions](#).

Related documentation

In addition to this document, please see other documents for this product.

For the following documentation, see the following web site:

<http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html>

- HP StorageWorks Continuous Access EVA overview
- HP StorageWorks Continuous Access EVA planning guide
- HP StorageWorks Continuous Access EVA Performance Estimator user guide
- HP StorageWorks EVA replication compatibility reference
- HP StorageWorks Replication Solutions Manager installation guide
- *HP StorageWorks Replication Solutions Manager online help and user guide*
- *HP StorageWorks Continuous Access EVA license key installation instructions*
- *HP StorageWorks Replication Solutions Manager administrator guide*

The following document is located on the HP StorageWorks JRE server CD:

- HP StorageWorks JRE Server installation guide

For information about the following products, see the respective web sites:

EVA3000

<http://h18006.www1.hp.com/products/storageworks/eva3000/index.html>

EVA4000

<http://www.hp.com/go/eva4000/>

EVA5000

<http://h18006.www1.hp.com/products/storageworks/enterprise/index.html>

EVA6000

<http://www.hp.com/go/eva6000/>

EVA8000

<http://www.hp.com/go/eva8000/>

Storage Management Appliances

<http://h18006.www1.hp.com/products/sanworks/managementappliance/index.html>

SAN design or SAN extensions

<http://h18006.www1.hp.com/products/storageworks/san/documentation.html>

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Blue text: Figure 1	Cross-reference links and email addresses
Blue, underlined text: http://www.hp.com	Web site addresses
Bold font	GUI elements that are clicked or selected, such as: <ul style="list-style-type: none">• Menu and list items• Buttons• Check boxes.
<i>Italics font</i>	Text emphasis
Monospace font	<ul style="list-style-type: none">• File and directory names• System output• Code• Text typed at the command-line
<i>Monospace, italic font</i>	<ul style="list-style-type: none">• Code variables• Command-line variables
Monospace, bold font	<ul style="list-style-type: none">• Emphasis of file and directory names• System output• Code• Text typed at the command-line



CAUTION:

Indicates that failure to follow directions could result in damage to equipment or data.

**NOTE:**

Provides additional information.

HP technical support

Telephone numbers for worldwide technical support are listed on the HP web site:
<http://www.hp.com/support/>.

Collect the following information before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

For continuous quality improvement, calls may be recorded or monitored.

HP strongly recommends that customers sign up online using the Subscriber's choice web site:
<http://www.hp.com/go/e-updates>.

- Subscribing to this service provides you with email updates on the latest product enhancements, newest versions of drivers, and firmware documentation updates as well as instant access to numerous other product resources.
- After signing up, you can quickly locate your products by selecting Business support, then Storage under Product Category.

HP-authorized reseller

For the name of your nearest HP-authorized reseller:

- In the United States, call 1-800-282-6672.
- Elsewhere, see the HP web site: <http://www.hp.com>. Click Contact HP to find locations and telephone numbers.

Helpful web sites

For other product information, see the following web sites:

- <http://www.hp.com>
- <http://www.hp.com/go/storage>
- <http://www.hp.com/support>
- <http://www.docs.hp.com>

Providing Feedback

We welcome your feedback!

For HP Command View EVA, please mail your comments and suggestions to CVFeedback@hp.com.

For HP Business Copy EVA and HP Continuous Access EVA, please mail your comments and suggestions to EVAReplication@hp.com.

1 About HP Continuous Access EVA

Overview

This chapter provides an overview of replication features, and gives a brief description of required hardware components and software applications.

Overview of HP Continuous Access EVA

HP StorageWorks Continuous Access EVA is the remote replication component of HP StorageWorks Enterprise Virtual Array (EVA) controller software. When this component is licensed, the controller copies data online, in real time, to a remote array over a local or extended storage area network (SAN). Properly configured, HP Continuous Access EVA is a disaster-tolerant storage solution that guarantees data integrity if an array or site fails.

Figure 1 shows a basic configuration with redundant arrays and fabrics. One array is located at a local (or active) site and the other at a remote (or standby) site. In the figure, one fabric is called the black fabric and the other is called the gray fabric. Each array can perform primary data processing functions as a source, with data replication occurring on the destination array. The replication process can also be bidirectional, with some I/O streams moving to the array and other I/O streams moving simultaneously from the array. This feature allows the array to be the source for some data groups and the destination for others. Figure 1 shows HP Continuous Access EVA in an EVA 3000/5000 configuration only.




Figure 1 Basic HP Continuous Access EVA configuration with redundant servers

Callouts:

1. Local active management server
2. Local host
3. Local controller 1

4. Local black fabric switch
5. Local controller 2
6. Local gray fabric switch
7. Local standby management server (optional)
8. Interswitch link–black fabric
9. Remote standby management server 1
10. Remote host
11. Remote controller 1
12. Remote black fabric switch
13. Remote controller 2
14. Remote gray fabric switch
15. Remote standby management server 2 (optional)
16. Interswitch link–gray fabric

In [Figure 1](#), the management server represents the server where EVA management software is installed, including HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager. HP recommends at least two management servers at each site to avoid a single point of failure. If a significant failure occurs at the source array location, redundancy allows data processing to quickly resume at the destination array. This process is called failover. When the cause of the array failure has been resolved, data processing can be moved back to the original source array by performing another failover.

Features

HP Continuous Access EVA's more prominent features include the following:

- Synchronous remote replication—ensures both source and destination copies are always identical and concurrent.
- Asynchronous remote replication—allows write completion back to the host for a shorter host I/O response time.
- Automated failover support.
- Normal and failsafe data protection modes of operation.
- Intersite link suspend—and—resume operations.
- Bidirectional replication enables two sites in a remote replication connection to use each other to maintain synchronized copies of online data.
- Source and destination pair size of 1 GB to 2 TB in 1 GB increments.
- Up to 128 remote source and destination pairs per array for HSV 110/100 controllers. Up to 256 remote source and destination pairs per array for HSV 210/200 controllers.
- Up to 8 source and destination pairs per DR group for HSV 110/100 controllers. Up to 32 source and destination pairs per DR group for HSV 210/200 controllers.
- HP StorageWorks Replication Solutions Manager host agents, enables you to mount volumes and run jobs directly on storage hosts.
- Job language and job templates to automate replication tasks.
- HP StorageWorks Replication Solutions Manager Command Line User Interface
- Replication is allowed between any two arrays. The maximum number of source and destination pairs and DR groups are limited to the smaller of those for each array.
- Auto suspend of DR replication activities if link between systems goes down.
- Improved write history logging support—EVA 4000/6000/8000 only. You now can choose the disk group where the system writes the log file.
- Internal job scheduler in RSM.
- May manually initiate a full copy control in EVA 4000/6000/8000 environments.

Setup and configuration assumptions

Before you can perform replication, the hardware and software components must be installed and a license activated. This guide assumes that you have the following components installed and configured.

Licensing

To perform remote replication, an HP StorageWorks array must have a valid HP Continuous Access EVA replication license (also known as a license-to-use, or LTU). To perform local replication—snapshots and snapclones—the array must have a Business Copy replication license. See the QuickSpecs for information on ordering a license. To activate licenses, see the documentation that comes with the license agreements.

Hardware and software components

The following sections describe the hardware and software components that must be installed before you can perform replication. See the *HP StorageWorks Continuous Access EVA planning guide* for more information.

Array


You must have at least two configured EVA arrays: one at a primary site and one at an alternate site a safe distance away.

Hosts and host operating systems

Use a supported operating system. See the *HP StorageWorks EVA replication compatibility reference* for a list of supported host operating systems and versions.

Virtual Controller Software

HP Continuous Access EVA is enabled in the Virtual Controller Software on redundant EVA arrays. The controller software provides the functionality for the array controller. [Figure 2](#) illustrates the role of the controller software in a single array. For additional information, see the HP StorageWorks Enterprise Virtual Array user guide.



CXO8166b

Figure 2 HP EVA management environment

Callouts:

1. Hosts
2. Fabric
3. Array controller software
4. Array
5. Administrator

6. Browser
7. Control input monitoring output
8. Control and monitor commands
9. Management server



NOTE:

Proper zoning ensures that hosts in an HP Continuous Access EVA environment do not conflict with each other. See the *HP StorageWorks Continuous Access EVA planning guide* for more information.

Management server and software

You must install and run HP Command View EVA on one of the following server types:

- Dedicated server—Runs only storage EVA software.
- Storage Management Appliance (SMA)—Runs software applications at a centralized location for managing and monitoring your storage. The SMA arrives preloaded with SMA software upon which other management applications can be loaded. HP StorageWorks Command View EVA and HP StorageWorks Element Manager for HSG software (preinstalled on the SMA), provide physical and logical views of your array through a graphical user interface.
- General purpose server—Runs other applications in addition to storage EVA software.



NOTE:

In this guide, the term management server applies to any of the server types above.

Interface options

Performing replication can involve one or more of the interfaces and products described in [Table 2](#).

Table 2 Other interfaces and products that perform replication

HP product	Interface	Remarks
HP StorageWorks Command View EVA	Browser-based graphical user interface (GUI)	<ul style="list-style-type: none">• Required.• Allows manual creation of snapshots and snapclones from a GUI.• Replicates by specifying an array and virtual disk; cannot replicate by specifying a host and host volume.• Cannot perform dynamic mounting or interactions with hosts.• Does not provide jobs, job templates, or scripting capabilities.• See HP Command View EVA for more information.

<p>HP StorageWorks Replication Solutions Manager</p>	<p>Browser-based GUI and command line user interface (CLUI)</p>	<ul style="list-style-type: none"> • Optional. • Specialized storage replication software that supports both local replication and remote replication. • Creates snapshots and snapclones using a GUI, jobs, and a CLUI. • Replicates by specifying an array and virtual disk (LUN) or by specifying a host and host volume. • Performs dynamic mounting and interactions with supported hosts. • Includes integrated job editor, job templates and job scripting capabilities. • Includes integrated job management.
--	---	--

HP StorageWorks Storage System Scripting Utility (SSSU)	Host command line	<ul style="list-style-type: none"> • Optional. • EVA host platform software that creates snapshots and snapclones from a command line or custom script. • See snapshot and copy commands in the EVA SSSU reference guide. • Replicates by specifying an array and virtual disk; cannot replicate by specifying host and host volume. • Dynamic mounting and host interactions is accomplished by writing custom scripts.
HP StorageWorks SMI-S Interface for Command View EVA	WEBM client-server using XML	<ul style="list-style-type: none"> • Optional. • Provides an SMI-S compliant interface for HP StorageWorks Command View EVA.

HP Command View EVA Tasks

HP Command View EVA is a user interface that communicates with the EVA controllers to control and monitor the storage. HP Command View EVA maintains a database for each managed array and the database resides on that array. Instructions for its use can be found in the HP StorageWorks Command View EVA Online Help.

To use HP Continuous Access EVA you must first use HP Command View EVA to:

- Add licenses
- Initialize controllers—This process binds the controllers together as an operational pair and establishes preliminary data structures on the disk array.
- Create disk groups—A disk group is a set of physical disks from which storage pools are created. When your array is initialized, one default disk group is created. For the highest performance and availability, a disk group should only contain one disk drive model. If you are performing bidirectional replication, the disk groups should be symmetric with respect to the capacity on both arrays.
- Create host definitions—The host connects to a fabric through a host bus adapter (HBA) and accesses storage through the controllers. Hosts contain a pair of HBA ports to connect with each fabric. Before a host can access any storage, it must be discovered by the array. This process assigns a host name to at least one HBA.
- Create virtual disks—Variable disk capacity that is defined and managed by the array controller and presentable to hosts as a disk. You can assign a combination of characteristics to a virtual disk, such as a name, redundancy level, size, and other performance characteristics.

- Present virtual disks to hosts—Assigning a virtual disk to a host results in a host presentation. You may present a virtual disk to a host during the virtual disk creation process, or wait until a later time. However, the virtual disk must be presented to a host in order to use it for replication.

You can also use HP Command View EVA to create snapshots and snapclones.

HP Replication Solutions Manager Tasks

HP Replication Solutions Manager is an interactive, visual environment for managing data replication. After you create hosts, virtual disks and their presentations, and your hosts can access your virtual disks, you are ready to replicate data using HP Replication Solutions Manager. Additional operational information is in the online help system.

HP recommends that you use the replication manager to:

- Create and delete DR groups.
- Change properties for DR groups.
- Add and remove members from DR groups.
- Failover, suspend, resume, and change failsafe mode by DR group.
- Create and delete managed sets.
- Failover, suspend, resume, enable and change failsafe mode, and revert to Home by managed set.
- Direct commands to the appropriate array regardless of selected array.
- Actively monitor arrays (copy, failsafe, logging, merging).
- Create and restore configuration databases for managed sets and jobs.

2 Concepts

Overview

This section describes some basic terminology and concepts you need to understand before replicating data.

Local replication

Local replication is a licensed feature of HP StorageWorks arrays that allows you to quickly create local, point-in-time copies of your data. These copies are known as snapshots and snapclones.

In a typical environment, a server is running HP Command View EVA and a local replication automation product, for example Replication Solutions Manager (hereafter called the replication manager).

The replication server is connected by LAN to multiple hosts running the replication manager host agents (hereafter called enabled hosts). These enabled hosts are running production and backup applications which perform I/O with multiple storage arrays in the SAN. An operator or administrator automates operations by running the replication manager jobs from a browsing computer, a scheduler, or from a host using the replication manager CLUI.

In this environment, the replication manager jobs can perform simultaneous, non-disruptive tape backups of the databases. To run daily backups, for example, you could create two jobs, each performing backups with their respective storage arrays and hosts.

You can also perform local replication using HP Command View EVA or any of the other supported HP Business Copy EVA interfaces, provided you have an HP Business Copy license.

Snapclone

Snapclone replication instantly creates a copy of a virtual disk that begins as a fully allocated snapshot, then becomes an independent virtual disk.

Snapshot

Snapshot replication instantly creates a demand allocated or fully allocated point-in-time copy of a source virtual disk. A demand allocated snapshot is a virtual copy in which the allocated disk space can change on demand from an initial minimum amount, up to the full capacity of the source. A fully allocated snapshot is a virtual copy in which the allocated disk space is initially set to, and remains fixed at, the full capacity of the source.

Remote data replication

The array at the active location is connected to a partner array at the standby location. To replicate data remotely, a source virtual disk is configured at the active array. When data replication is selected, the destination virtual disk (called a remote copy) is automatically created by software at the standby array. Any data written to the source virtual disk is then copied to the destination virtual disk. Applications continue to run while data replicates in the background over a separate connection. The direction of remote data replication is set as a property of the DR group.

Bidirectional replication

When an array contains both source virtual disks and destination virtual disks, it is said to be bidirectional. An array can have a bidirectional data replication relationship with up to two other arrays; and an individual virtual disk can have a unidirectional replicating relationship with only one other virtual disk.

DR groups

A data replication (DR) group is a named group of virtual disks selected from one or more disk groups so that they remotely replicate to the same destination, fail over together, share a log (DR group log disk), and preserve write order within the group.



NOTE:

To use this feature with specific source and destination arrays, each array must have its own HP Continuous Access EVA replication license. Also note:

- Only virtual disks eligible for remote replication can belong to a DR group.
 - A DR group can contain one or more virtual disks. All vdisks assigned to an application must be in the same DR group.
 - You can create up to 128 DR groups per storage array for EVA 3000/5000 controllers and 256 DR groups for EVA 4000/6000/8000 controllers.
-

Source and destination DR groups

DR groups are always created in pairs, consisting of a source DR group and a destination DR group. The source DR group contains the virtual disk you want to replicate remotely. This is the source virtual disk. The destination virtual disk (called the remote copy) resides in the destination DR group.

Replication direction

The replication direction of a DR group is always from a source to a destination. When replicating from a source to a destination, the DR Group is in an original state. When replication occurs from an array that was created as a destination (for example, after a failover), the DR group is failed over.

Figure 3 depicts the replication of one DR group between separate sites.




Figure 3 HP Continuous Access EVA DR group replication

Callouts:

1. Host
2. Switch
3. Host I/O
4. Replication writes
5. Array A
6. Array B
7. Source virtual disk
8. Destination virtual disk
9. DR group containing a pair of source and destination virtual disks

Home designation

Home is the preferred source DR group in a remote replication relationship. Having a preferred source allows you to designate a point of reference for remote replication. Home designation is set in the RSM GUI. Home designation is not seen in HP Command View EVA.

DR group presentation

All members of a DR group are presented to a host through the same controller and move together. For optimal performance, limit each DR group to one virtual disk.



NOTE:

Use the same preferred path for all DR group members, with presentation to same host bus adapters (HBAs).

Presenting DR group members to the same HBA

For EVA 3000/5000 controllers, all members of a DR group must be presented to the same HBA on hosts with more than one HBA per fabric (for example, multiple HBA pairs, multiple dual-channel HBAs, or a combination of single and dual-channel HBAs). All DR group members must also be preferred to the same controller with the same failover characteristics.

This restriction is required to keep the DR group members using the same host HBA to EVA path. In the event of a path or controller failure, the members collectively fail over to the other path, thus preserving write order across the members of the DR group. You must turn off dynamic load balancing in SecurePath.



NOTE:

Additional members added to a DR group will have the parameters of the original member. This can affect multipathing of OS applications. For the EVA 4000/6000/8000, you must have a single active SCSI initiator per vdisk. You can use any array port to access any vdisk.

DR group properties

Every DR group has defined properties such as name, operational state, and so on. See the online help for an explanation of the properties.

DR group log

The DR group log is a designated virtual disk that stores a source DR group's host writes while remote replication to the destination DR group is stopped or if the writes to the source side are faster than transfer rate to the destination side of the DR group. When replication is re-established, the contents of the log are written to the destination virtual disks within the destination DR group to synchronize the destinations with their sources. This process of writing the log disk contents to the destination in the order that the writes occurred is called merging.

Once a log reaches full capacity, it is often more practical to copy the changed blocks on the source virtual disk to the destination virtual disk. This copy operation is called a fast synchronization—all 8-MB block increments written on a source virtual disk since it lost connection with the destination are copied to the destination virtual disk.

At other times, a full copy copies the complete source virtual disk to the destination virtual disk. The full copy is an automatic process that occurs when a log has reached full capacity.

The force full copy feature is only allowed in a HSV 210/200 controller source *and* destination environment.

DR group log states

A DR group log can be in one of the following states:

- Unused (Normal)–No source virtual disk is logging or merging.
- Logging–At least one source virtual disk is logging but none are merging.
- Merging–At least one source virtual disk is merging and logging.

DR group log size

When a DR group is created, a log disk consists of 136 MB of Vraid1 space; however, the array can automatically allocate more space on demand. When a DR group is logging, the log disk grows in proportion to the amount of writes to the source virtual disk.

When creating disk groups and the virtual disks within them, ensure that sufficient space remains for DR group log disks to expand to their maximum size. HP recommends creating DR group log disks in near-online disk groups, if available. Otherwise, create log disks in the online disk groups with the most free space.



NOTE:

You can select the disk group for the log on an EVA 4000/6000/8000 system; the array controller software will select it for you for EVA 3000/5000 systems. See the *HP StorageWorks Continuous Access eva planning guide* for more information.

Array software considers the log disk full when any of the following conditions occurs:

- No free space remains in the disk group.
- The log disk reaches 2 TB of Vraid1 (4 TB total).
- The log reaches the maximum log size, due to either reaching the user-set maximum or twice the DR group's virtual disk combined size.

When the log disk is declared full, DR group members are marked for full copy and the log disk is deleted.

Managed sets

A managed set is a named collection of resources banded together for the purpose of management. For example, the managed set Sales_Disks might include two virtual disks, West_Sales and East_Sales.

Performing an action on a managed set, performs the action on all the members in the set. For example, if you perform the New Snapshot action on the managed set Sales_Disks, the interface creates a snapshot of West_Sales and a snapshot of East_Sales.



NOTE:

The order in which members are added to the group has no effect on the order in which actions are performed on the members.

Members

A managed set can comprise the following types of resource: DR groups, enabled hosts, host volumes, storage systems, or virtual disks.

Managed set properties

A managed set is defined by its properties, such as name and type. Use the replication manager to view the properties.

Managed set actions

Each type of managed set supports common actions, plus resource-specific actions (see List tab actions and Members tab actions). Use the Managed Sets List and Tree tabs to view the managed sets and members managed by the replication manager, and to perform actions.

Considerations

General considerations

- All resources, or members, in a single managed set must be of the same type (for example, all virtual disks).
- You can add a specific resource to more than one managed set.

Storage resource considerations

You can add resources on more than one storage system to a managed set.

DR group considerations

- Source and destination sides of different DR groups can be part of the same managed set. However, both the source and destination sides of the same DR Group cannot be part of the same managed set.
- Some DR group actions are permitted only on source DR groups; other actions are only permitted on destination DR groups. See the DR group actions table for more information.
- Create separate managed sets for source and destination DR groups so that if a failover occurs, you can perform the actions that correspond to the new identity of the managed set.
- If you plan to use DR group managed sets for failover operations, ensure the managed sets are controlled by the same server at the time of failover.

Failover

Failover is an operation to reverse replication direction. When you initially set up data replication, you replicate from a source to a destination. Failover simply changes the direction of replication; the destination array assumes the role of the source and the source assumes the role of the destination. For example, if a DR group is replicating from array Alpha to array Bravo, a failover operation would change the direction of the replication so that data from array Bravo would be replicated to array Alpha.


You can fail over:

- DR groups—Reverses the replication direction of a DR group.
- Managed sets containing DR Groups—Reverses the replication direction of all DR groups within the managed set.

Figure 4 shows a data replication relationship among DR groups at three locations. Arrays A and D are located at the primary site and arrays B and C are located at two alternate sites. If contact with the primary site is broken, failover can manually occur with the destination DR groups that were replicating to array B. Array B then acts as the primary site for these DR groups after failover. At array C, the source DR groups begin logging until new remote sites are re-established.

Considerations

- Failing over is only permitted on a destination DR group.
- If only one component fails, repairing that single component may be preferable to performing a complete site failover.
- Consider using the using the auto suspend on link down feature to choose how to work through a backlog of jobs after a failure.



CXO7993c

Figure 4 Replicating relationships among DR groups

Callouts:

1. Source array A
2. Destination array B
3. Replication
4. Destination array D

5. Source array C
6. Active site
7. Standby site
8. Failover
9. Logging

Replication manager logs

The following logs track and record activity in the replication manager:

- Event log—Contains system-generated messages resulting from:
 - User-initiated actions (for example, "suspend DR group").
 - Replication-related storage events (for example, "DR group is constructing").
 - Jobs (for example, "job complete").

Messages are written to a set of five rotating files. The default size limit for each of these files is 1 MB. As the current file reaches its size limit, it is closed, rotated out, and a new file opened. You can view the messages in the Event pane of the replication manager. See the online help for more information.

- Trace log—Contains all events. Intended for HP personnel. Events that are useful to the user are transferred to the Event log and displayed in the Event pane. The Trace log is allotted 60 MB of space. As the log fills the space, old messages are discarded. You can view the Trace log in the Configuration window in the replication manager.
- Security log—Contains the following security activities:
 - Successful attempts to access the replication manager
 - Failed attempts to access the replication manager
 - Failed authorization credentials
 - Changes to security user accounts and group membership

Thirty days of history are saved by default. On the thirty-first day, the oldest entries are discarded. You can change the default to any number of days. HP recommends you save no more than 30 days. The file to change is called `security.cfg` and is located on the replication manager server in `C:\Program Files\Hewlett-Packard\sanmgr\security\config\`.

Security logs are .txt files located on the replication manager server in `C:\Program Files\Hewlett-Packard\sanmgr\security\logs`.

- Transaction log—Contains the replication manager database contents. Service personnel can access this log to recover data if a database is corrupted. No user action is needed or allowed on this log. Space needs vary for the log depending on database activity. However, space needs are generally less than 20 MB.

Failsafe mode

The failsafe mode specifies how host writes and remote replication behave when a group member fails. The failsafe mode can be either:

- Failsafe enabled—If any virtual disk within the DR group fails or becomes unreachable, all host writes and remote replication automatically stop. This preserves the order of the replicated data. A failsafe-enabled DR group can be in one of two states:
 - Locked (failsafe-locked)—Host writes and remote replication automatically stop.
 - Unlocked (failsafe-unlocked)—Host writes and remote replication occur.
- Failsafe disabled—If any destination virtual disk (remote copy) within the DR group fails or becomes unreachable, all host writes to the source DR group continue, but all remote replication to the destination DR group automatically stops; the source DR group logs its host writes to the DR group log until remote replication is re-established. If a source virtual disk fails, host writes to the failed disk stop, as well as remote replication to its remote copy; host writes and remote replication to the other members of the DR group continue normally.

3 Recovery

Overview

This chapter provides recovery information for performing failovers, and resuming operations after encountering a failsafe-locked condition or a disk group failure. Procedures use the HP StorageWorks Replication Solutions Manager, when applicable. Several scenarios are provided that cover most situations you could encounter, with procedures for handling each scenario.

Planning for a disaster

When a disaster occurs at one of your storage sites, your first priority is to get your data back online in the shortest amount of time. Planning helps to minimize the downtime brought on by a disaster:

- Operating with a supported disaster-tolerant HP Continuous Access EVA configuration is of primary importance. Ensure that there are two fabrics with at least one intersite link per fabric.
- Ensure that your controllers are cabled in the supported configuration to your fabrics. See the *HP StorageWorks Continuous Access EVA planning guide* for more information.
- Have at least one management server available at every site, in case of a hardware or communication failure.
- Verify that each destination virtual disk within a DR group has been presented to a host. This allows the host access to the virtual disk immediately after a failover.
- Ensure that local and remote hosts are installed with the latest patches, virus protection, EVA platform kits, and Secure Path versions for that operating system.
- Keep your configuration current and documented at all sites. Install the latest versions of the virtual controller software firmware, management server software, HP Command View EVA, and the replication manager.
- Keep a record of your virtual disks and DR groups. Capture the configuration information after significant changes or at scheduled intervals (see [Backing up configuration information](#)).
- Keep the replication manager on every management server synchronized with any configuration changes.
- Back up the replication manager database for any configuration that the management server may use. These databases contain managed set and job information that you can quickly restore when a management server changes its role. For example, two management servers can control different arrays at one time, but each management server can have a database created with all the arrays under its control in case the other one becomes inoperative.
- Practice the recovery plan. Ensure that everyone involved in your storage administration practices for disaster recovery. Practice different failure scenarios and make decisions ahead of time about them. For example, if a controller fails, is it more important not to disrupt processing by doing a planned failover, or to not be at risk for a second controller failure that will result in an unplanned failover? In the case of multiple sites, which site has precedence for troubleshooting?

Scheduling practice disaster recoveries is a good way to verify that your records are up-to-date and that all required patches are installed.

Failover

Failover can take several forms with HP Continuous Access EVA:

- Controller failover is the process that takes place when one controller in a pair assumes the workload of a failed or redirected companion controller in the same cabinet.
- DR group or managed set failover is an operation to reverse the replicating direction of the DR group or managed set to its partner. This reversal is possible because all data generated at a source array has been replicated to a destination array, in readiness for such a situation.
- Fabric or path failover is the act of transferring I/O operations from one fabric or path to another.

This chapter discusses the failover of DR groups and managed sets. It does not discuss controller failover within a cabinet, or path or fabric failover, because redundancy is assumed.

The failover method used with DR groups, managed sets, or arrays is determined by the severity of the failure or the reason for the failover. A planned failover can be used for situations such as an anticipated power disruption, scheduled equipment maintenance at the source array, or the need to transfer operations to another array. An unplanned failover is used for events such as multiple controller failures, multiple host failures, or an unplanned power outage at the source array.

If the source array fails, or if you are planning downtime with the source array, you must decide whether to perform a failover to the destination array. Always verify that all components at the destination array are operational before you begin a failover.

When you perform a failover, the destination array assumes the role of the source and becomes the active array. It remains the source array until you fail over to another system. By transferring control of system operation to the destination array, you can ensure minimal interruption of data access after a failure.



NOTE:

When you perform a failover for a DR group or a managed set, you must fail over all components of the group or set. Therefore, if only one component fails, repairing that single component may be preferable to performing a complete failover. HP recommends that you not perform a planned or unplanned failover of one or more DR groups more frequently than once every 15 minutes. The planned or unplanned failover of a controller should also not be performed more frequently than once every 15 minutes.

[Table 3](#) outlines example situations that require a failover and those that do not. For each type of failover, an action is recommended, which may require action at the source or destination array. Since replication

can be bidirectional, one array can be the source and destination for different DR groups. You can use this table to customize contingency plans within your specific environment.

Table 3 When to and when not to fail over a DR group, managed set, or array

Type of failure	Recommended action	
	DR group in normal mode	DR group in failsafe mode
Total loss of source array	Manually intervene to fail over data to destination array, then restart processing at the destination array (see Unplanned failover).	
Loss of both source controllers		
Loss of single source controller	Failover not necessary.	
Total destination array loss	Failover not necessary.	Manually intervene to continue processing at source array. See Resume operations if unable to access destination .
Loss of both destination controllers		
Loss of all intersite links		
Loss of both source intersite switches	Manually intervene to fail over data to destination array, then restart processing at the destination array (see Unplanned failover).	
Loss of single source intersite switch	Failover not necessary.	
Extended power outage at primary site	Manually intervene to fail over data to destination array, then restart processing at the destination array (see Unplanned failover).	
Loss of a managing server	Failover not necessary. Browse to standby managing server.	
Loss of single disk in redundant storage	Failover not necessary.	
Loss of single host of cluster	Failover not necessary.	
Disk group hardware failure (loss of redundancy) on the source array	Fail over to destination, and repair array (see Disk group hardware failure on the source array).	
Disk group hardware failure (loss of redundancy) on the destination array	Failover not necessary (see Disk group hardware failure on the destination array).	

Backing up configuration information

It is important to have a record of your storage configuration in case of a hardware failure (see [Return operations to replaced new storage hardware](#)).

Using SSSU to capture configuration information

One way to capture your configuration information is with the Storage System Scripting Utility (SSSU).



NOTE:

The SSSU is not available on Novell NetWare hosts or on the management server.

Using the SSSU `Capture Configuration` command, five scripts are run that append a user-defined configuration name to the file with a name in the form `UserName_StepX.txt`, where StepX is one of these five configuration text files:

- Step1A–Captures the data needed to re–create the array itself, disk groups, hosts, virtual disks that are not used for data replication (either source or destination), and LUNS for the disks created.
- Step1B–Captures the data needed to re–create all source virtual disks used in DR groups on this array. However, the data captured by this step is not currently used in any recovery procedures.
- Step1C–Captures the data needed to present all source virtual disks (creates LUNs) used for DR groups to their hosts. However, the data captured by this step is not currently used in any recovery procedures.
- Step2–Captures the data needed to re–create all data replication specific configuration information only, and only DR–specific information for which the array is the source. This consists of source DR groups and their members only.
- Step3–Captures the data needed to create an SSSU script that will again present all remote virtual disks to their hosts.

Example:

You have a management server with an IP address of 111.222.333.444. You want to back up your configuration for an array named HSV01 with the SSSU. Follow these steps:

1. Run the SSSU executable (SSSU.exe) to get a command prompt.
2. At the command prompt, log onto the management server using your user name and password. For example:

```
select manager 111.222.333.444 username=user1 password=admin
```

3. Select the array whose configuration you want to save. Use the following command with your array name (HSV01 is the array name used in this example):

```
select system HSV01
```

The command line prompt changes to reflect your array is selected.

4. At the SSSU command prompt, enter the `Capture Configuration` command along with a path and file name where you want the configuration text files to reside. For example, to copy these files to a folder called `storage_systems\hsv01`, use the command:

```
capture configuration c:\storage_systems\hsv01.txt
```

Messages on the screen confirm that each step was successfully saved.

A current copy of the data that defines the array configuration is saved. If the array configuration changes later, rerun this procedure. Procedures to recover a configuration using the SSSU are in [Return operations to replaced new storage hardware](#).

Manually capturing configuration information

You can capture configuration information by writing it down manually. Use the following form as a guideline for capturing the information.

Table 4 Manual configuration form

Array name:	
Array WWID:	
Console LUN ID:	(default = 0)
Disk group information	
Disk group name:	(default = default disk group)

Device count:	
Spare policy:	(none, single, or double)
Disk type:	(online or nearline)
Occupancy alarm:	(default = 95%)
Host information	
Folder name:	(default = \Hosts\)
Host name:	
Operating system:	
For each HBA port:	WWID:
Virtual disk information	
Folder name:	(default = \Virtual Disks\)
Virtual disk name:	
Disk group:	
Size:	
Redundancy level:	(Vraid0, Vraid1, Vraid5)
Write cache policy:	(Mirrored write-back, unmirrored write-back)
Read cache policy:	(On, off)
Read/write permissions:	(Read/write, read-only)
OS unit ID:	(default = 0)
Preferred path:	(None, Path A FO, Path A FO/FB, Path B FO, Path B FO/FB)
Presentation:	Host name:
	LUN:
DR group information	
Source:	Array name:
	Virtual disk members:
Destination:	Array name:
	Virtual disk members:
Parameters:	Failsafe mode: (disabled, enabled)
	Write mode: (synchronous, asynchronous)
	Destination mode: (none, read-only)

Possible event scenarios

This section discusses the following scenarios that require manual intervention:

- [Planned failover](#)
- [Unplanned failover](#)
- [Resume operations if unable to access destination](#)
- [Return operations to Home array](#)
- [Return operations to replaced new storage hardware](#)
- [Disk group hardware failure on the source array](#)
- [Disk group hardware failure on the destination array](#)

Performing recovery actions

Recovery procedures require such actions as failover, suspend, resume, disable failsafe, mounting, and unmounting. You can perform these actions using various interfaces and tools:

- The replication manager
- Command line user interface
- Job scripting
- Storage System Scripting Utility
- HP Command View EVA

This chapter describes generic recovery procedures. For specific procedures, see the documentation for your preferred method.

Planned failover

Situation: Due to scheduled maintenance at the primary site, you need to perform a planned move of operations from the source array to the destination array.

Action: Prepare the source array for the failover, then perform a failover to the destination array. After the failover is complete, you can continue to operate from this array and revert back to failsafe mode, if desired. When the maintenance is complete you can failover to the original source array (see the procedure [Planned failover](#)).



NOTE:

You can set the Home designation in the replication manager interface to identify the preferred array. By default, an array created as a source is designated as the Home array. Because the role of the source array can change during failover, this Home designation allows you to identify your preferred array.

Unplanned failover

Situation: You have experienced an unplanned loss of the primary site or the source array. The duration of the outage at the source is unknown. The HP Continuous Access EVA hardware components (hosts, controllers, and switches, for example) at the primary site may or may not remain intact.

Action: Perform an immediate failover to the remote site or passive array. When the primary site is back online, you can choose to return to the Home array or to one with new hardware (see the procedure [Unplanned failover](#)).

Resume operations if unable to access destination while source in failsafe-locked state (extended period of time)

Situation: You have experienced an unplanned loss of the destination array, or a loss of the connection to the destination array, due to failure of the intersite links, loss of power at the alternate site, loss of both destination switches, and so on. The duration of the outage is unknown. The DR groups are in failsafe-enabled mode and host I/O is paused because the DR groups are failsafe-locked.

Action: Change from failsafe-enabled to normal mode, then resume host I/O until the connection to the destination array is re-established. When the connection to the destination site is stable, change back to the failsafe-enabled mode (see the procedure [Resume operations if unable to access destination](#)).

Return operations to Home array

Situation: You are operating from an array that was not originally designated as Home within the replication manager. You need to perform a planned move of operations from this alternate source array to the Home array.

Action: Prepare the Home array for the failover, then perform a failover to the Home array (see the procedure [Revert to Home \(failback\)](#)).

Return operations to replaced new storage hardware

Situation: Some type of disaster (lightning, flood, fire, severe equipment failure, or so on) has damaged equipment at the Home site and forced a failover to an alternate site. You are operating from an array that was not originally designated as Home.

Action: When the damaged components at the Home site (hosts, controllers, or switches, for example) have been repaired, and the site is operational and back online, perform a failover to new hardware at the Home site (see the procedure [Return operations to replaced new storage hardware](#)).

Disk group hardware failure on the source array

Situation: A hardware failure on your source array causes a disk group to become inoperative. This can be caused by the loss of enough disks to create a loss of redundancy within the disk group and affects all Vraid types present on the disk group.

Action: If you plan to recover using data on the destination array, then failover to the destination array. Delete DR groups and virtual disks on the failed array. Repair the failed disk group. Re-create DR groups, virtual disks, and host presentations.

If the failed source array was logging at the time of the hardware failure, you must recover with data at the destination site or from a backup.

Disk group hardware failure on the destination array

Situation: A hardware failure on your destination array causes a disk group to become inoperative. This can be caused by the loss of enough disks to create a loss of redundancy within the disk group and affects all Vraid types present on the disk group.

Action: Delete the DR groups on the source array that replicated to the failed disk group. Repair the failed disk group on the destination array. Re-create your DR groups on the source array and make host presentations at the destination array.

Failover and recovery procedures

The following procedures explain how to resolve the scenarios mentioned previously. HP recommends that you practice these procedures so that you are prepared in a crisis. Customize these procedures for your own use, if needed.

Planned failover

See [Figure 5](#) for the flow of steps required for a planned transfer of operations to a remote site. Complete the following steps:

1. If desired, move storage management to another management server.
2. All DR groups must be resumed. Check to ensure that full normalization has occurred. If a merge and full copy are occurring, wait for them to complete.
3. Stop all host I/O on the source array. Follow the steps listed in [Stopping host I/O](#).
4. Perform the failover operation.
5. If you plan to throttle I/O to specific arrays, suspend your less important DR groups at your new source. This forces the controllers to replicate the most important data first when the links to the previous source controller are re-established. See [Throttling of merge I/O after logging](#).

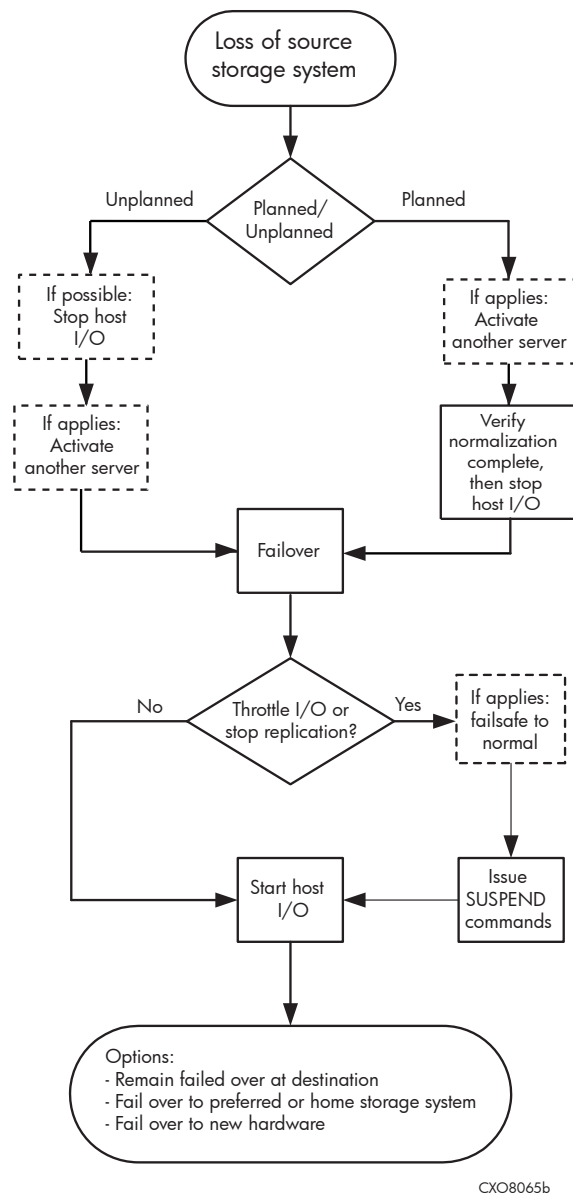


Figure 5 Planned and unplanned transfer of operations

6. If you plan to operate for an extended time at the alternate site (Home array and Fibre Channel links must be functioning properly) and you have a DR group that needs failsafe mode enabled, perform these steps:
 - If DR groups were suspended, resume copying on affected destination DR groups. Wait for the log disk to finish merging.
 - Change affected DR groups to failsafe mode.



NOTE:

You can enable failsafe mode at the destination array while a merge or full copy is being performed.

7. Issue operating system–dependent commands for presentation of units to remote hosts to start host I/O. Refer to [Allowing remote host discovery of devices](#).

After the transfer of operation is complete, you have three options after the cause of the failover is resolved:

- Remain failed over at the alternate (or destination) site.
- Return operations to the Home array (see [Revert to Home \(failback\)](#)).
- Return operations to new hardware (see [Return operations to replaced new storage hardware](#)).

Unplanned failover

See [Figure 5](#) for the flow of steps required for an unplanned transfer of operations to a remote site. Complete the following steps:

1. If your hosts are running on the source array, and you are able to access these hosts, stop all host I/O, as described in [Stopping host I/O](#).
2. If you cannot access the management server managing the arrays, establish management control with another management server.
3. Perform a failover to the destination site.
4. If you plan to throttle I/O to specific arrays, suspend your less important DR groups at your new source. This forces the controllers to replicate the most important data first when the links to the previous source controller are re-established. See [Throttling of merge I/O after logging](#).
5. Issue operating system–dependent commands for presentation of units to remote hosts to start host I/O. Refer to [Allowing remote host discovery of devices](#).

After the transfer of operation is complete, you have three options after the cause of the failover is resolved:

- Remain failed over at the alternate (or destination) site.
- Return operations to the Home array (see [Revert to Home \(failback\)](#)).
- Return operations to new hardware (see [Return operations to replaced new storage hardware](#)).

Resumption of operations if unable to access destination while source in failsafe–locked state (extended period of time)

See [Figure 6](#) for the flow of steps required to resume operations if you are unable to access the destination while in a failsafe–locked state. Complete the following steps:

1. Change affected source DR groups from failsafe–enabled mode to normal mode.
2. If necessary, issue operating system–dependent commands to the local hosts to start I/O again on the units that were failsafe–locked.
3. If you plan to throttle I/O to specific arrays, suspend your less important DR groups. This forces the controllers to replicate the most important data first when the links are re-established. When ready to merge to destination from source, issue the `Resume` command. See [Throttling of merge I/O after logging](#).

**NOTE:**

If you stay in a suspended state for an extended length of time, you can overrun the log, which initiates a full copy. During a full copy, the data is not usable until the full copy completes.

4. When connections to the destination site are re-established and merging is complete, change DR groups from normal mode to failsafe-enabled mode, if desired.

**NOTE:**

If source DR groups go into full copy mode, you can also enable failsafe mode.

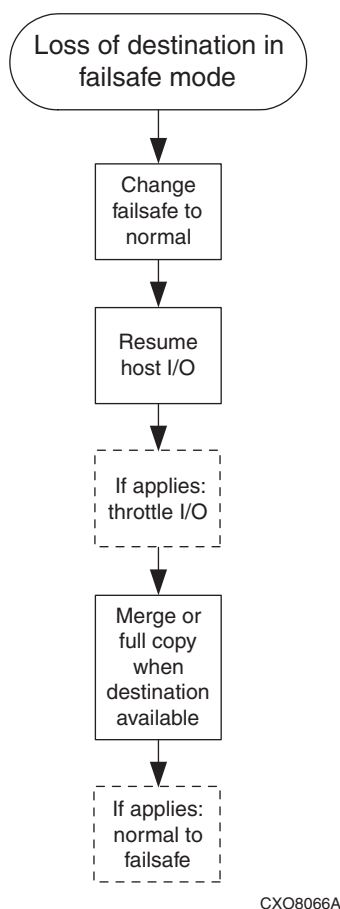


Figure 6 Resumption of operation if unable to access destination in failsafe mode

Revert to Home (failback)

Failback (also known as reverting to Home) is similar to a planned failover. After performing a failover from a source to a destination array, a failback can be performed in the other direction after waiting a minimum of 15 minutes. Complete the following steps:

1. If desired, move storage management to another management server.

2. All DR groups must be resumed. Ensure that full normalization occurred. If a merge and full copy are occurring, wait for them to complete.
3. Stop all host I/O on the source array. Follow the steps listed in [Stopping host I/O](#).
4. Perform the "revert to Home" operation. Or, perform the failover operation on affected DR groups.
5. If you plan to throttle I/O to specific arrays, suspend your less important DR groups at your new source. This forces the controllers to replicate the most important data first when the links to the previous source controller are re-established. See [Throttling of merge I/O after logging](#).
6. If you plan to operate for an extended time at the alternate site (Home array and Fibre Channel links must be functioning properly) and you have a DR group that needs failsafe mode enabled, perform the following steps:
 - If DR groups were suspended, resume copying on affected destination DR groups. Wait for the log disk to finish merging.
 - Enable failsafe mode of the affected DR groups.



NOTE:

You can enable failsafe mode at the destination array while a merge or full copy is being performed.

7. Issue operating system-dependent commands for presentation of units to remote hosts to start host I/O. Refer to [Allowing remote host discovery of devices](#).

After the transfer of operation is complete, you have three options after the cause of the failover is resolved:

- Remain failed over at the alternate (or destination) site.
- Return operations to the Home array (see [Revert to Home \(failback\)](#)).
- Return operations to new hardware (see [Return operations to replaced new storage hardware](#)).

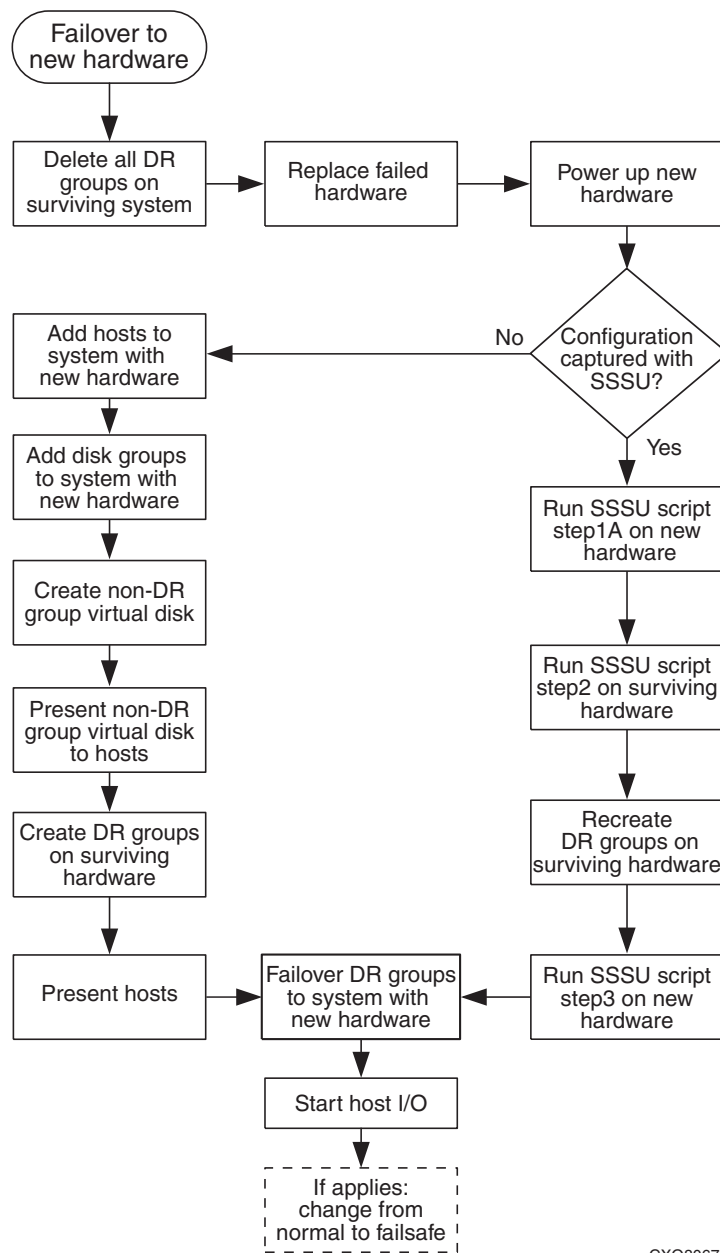
Return operations to replaced new storage hardware

This procedure is used after a failure that results in the replacement of array hardware at what was the source array. The procedure does not include steps to rebuild servers using the storage (this should be part of your overall disaster plan). The new hardware now acts as the destination array after a failover and is referred to in this procedure as the system with new hardware, or the array with failed hardware. The surviving system is now your source array after the failover. The steps below explain the process to return operations to a system having replaced new hardware:

Table 5 Array log

	Array with failed or new hardware	Surviving array
Array Name		
Array Name		
Array Name		
Array Name		
Array Name		
Array Name		
Array Name		
Array Name		

1. Denote your array names having failed or new hardware (destination) and your surviving array (source) in the table provided above. For example, your array with new hardware may be named HSV01 and your surviving array may be named HSV02. See the table during the procedure as needed.



CXO8067B

Figure 7 Array names having failed or new hardware (destination) and your surviving array (source)

2. Delete all DR groups on the surviving system that ever had a relationship with the failed hardware.
3. Replace the failed hardware. Depending on the failure, this means replacing hard drives or controllers, deleting disk groups, and so on.
4. Remove the connection between the source and destination arrays. This can be accomplished by removing it from the SAN, disabling the intersite links, or by placing the arrays into separate zones.

To place the array into separate zones, you need two zones. One zone contains the source array, source hosts, and the management server. The second zone contains the destination arrays, destination hosts, and the management server.

5. (Optional. Not needed if entire array was replaced.) Delete any destination DR groups on the previously failed array. If this is not successful, the source–destination connection still exists, so go to the previous step.
6. (Optional. Not needed if entire array was replaced.) Delete all virtual disks that were members of DR groups on the destination array.
7. Re-establish communication between the source and destination arrays. Either add the array back into the SAN, enable the intersite links, or place the arrays into the same zone.
8. Perform one of the following:
 - If the replaced array configuration was captured with the Storage System Scripting Utility (SSSU), execute the script `ConfigName_step1A` on the new hardware, then proceed to [Step 13](#). See the SSSU documentation for instructions. `ConfigName` is a user–assigned name given to the SSSU script at the time of creation. See the procedure titled [Backing up configuration information](#).
 - If you are not using an SSSU script for recovery, initialize the newly replaced array using the information you recorded in the [Table 5](#). See the HP Command View EVA documentation for initialization instructions.



NOTE:

To preserve your existing zoning, give the new hardware the same World Wide Names as they existed with the failed hardware.

9. Add the disk groups on the new hardware.
10. Add the hosts for the system with new hardware.
11. Create the non–DR group virtual disks.
12. Present all non–DR group virtual disks to their hosts.
13. Perform one of the following:
 - If the surviving array configuration was captured with the SSSU, execute `ConfigName_step2` on the surviving array. `ConfigName` is a user–assigned name given to the SSSU script at the time of creation. DR groups are re-created with the SSSU if they were performing as the source when the configuration was captured. This step may take some time to complete.
 - If you are not using an SSSU script for recovery, re-create all DR groups on the surviving system with destinations set to the new system using the information you recorded in the [Table 4](#).
14. If, in the previous step, you used the SSSU to re-create DR groups on the surviving array, you must manually re-create any additional DR groups that had their source on the failed hardware on the surviving array. This is necessary because the SSSU will not re-create those DR groups on the surviving array if they performed as the destination when the configuration was captured. After you perform this step, all DR groups reside on the surviving array.
15. At this point, you have the option of setting all affected DR groups from normal mode to failsafe-enabled mode.
16. Perform one of the following:

- If the original array configuration was captured with the SSSU, then execute *ConfigName_step3* on the new hardware. *ConfigName* is a user-assigned name given to the SSSU script at the time of creation.
 - If you are not using an SSSU script for recovery, present the destination virtual disks on the system with new hardware to the appropriate hosts using the information you recorded in the [Table 4](#).
17. If, in the previous step, you used the SSSU to present destination virtual disks to their hosts, manually present any additional virtual disks that originally had their sources on the failed hardware to their hosts on the array with new hardware. This is necessary because the SSSU will not present virtual disks whose destinations were the surviving array when the configuration was captured. After performing this step, all destination virtual disks are presented to hosts.
 18. If the system is to be the source for the DR groups, fail over any DR groups to the new array using the procedure [Planned failover](#).
 19. Issue operating system-dependent commands for presentation of units to remote hosts to start host I/O. Refer to [Allowing remote host discovery of devices](#).
 20. (Optional) Set the DR groups to the desired Home setting.

Recovering from a disk group hardware failure

Disk group hardware failure occurs when a disk group loses a quantity of disks beyond the capability from which a given Vraid type can recover. It is a loss of redundancy that results in an inoperative disk group. This condition can occur from the loss of one disk for Vraid0 to as few as two disks for Vraid1 and Vraid5. In each case, the hardware failure needs to be fixed, and the disk group data has to be structurally rebuilt. This section describes the symptoms and recovery methods for an inoperative disk group at either the source or destination array.




If an array only has one disk group, and that disk group fails, the array becomes inoperative. Re-initialize the array to manage it (see [Return operations to replaced new storage hardware](#)).

If you have multiple disk groups and one fails, follow the procedures on [Disk group hardware failure on the source array](#) and [Disk group hardware failure on the destination array](#).

Failed disk group hardware indicators

If disk group hardware fails, the replication manager displays the following:

Table 6 Replication manager display icons

Resource	Symbol	Description
Array		Indicates the array is in an abnormal state and requires attention.
Virtual disks		Indicates a catastrophic failure and requires immediate action.
DR groups		Red indicates a failure; yellow indicates the DR group is in a degraded state. Either condition requires immediate attention.

Disk group hardware failure on the source array

There are two ways to recover from a disk group hardware failure on the source array:

- If data replication was occurring normally when the source disk group became inoperative, the data at the destination array is current. A failover is performed to the destination array, DR groups are deleted, the inoperative disk group is repaired, and the DR groups are re-created. Data is then copied back.
- If your disk group becomes inoperative when your DR groups are logging (for example, your DR groups were suspended, or the intersite links are down), your data is stale on the destination array. Stale data is older data that is not as current as what exists on its partnered array. If you prefer to use stale data for recovery, the steps are the same as if replication was occurring normally. However, if you prefer to continue from a point-in-time, copy and then repair the inoperative disk group, and data is restored from a backup or full copy.



NOTE:

When you delete DR groups to recover from a disk group hardware failure, you lose the redundancy of the other site or disaster tolerance of your data.

Perform this procedure using HP Command View EVA when a disk group hardware failure occurs on the source array and the data on the destination array is current.

1. Navigate to each DR group on the surviving array and perform a failover (see [Unplanned failover](#)).
2. In HP Command View EVA, begin troubleshooting the disk group problem.
3. Navigate to the failed disk group.

A list of failed virtual disks and DR groups is displayed.

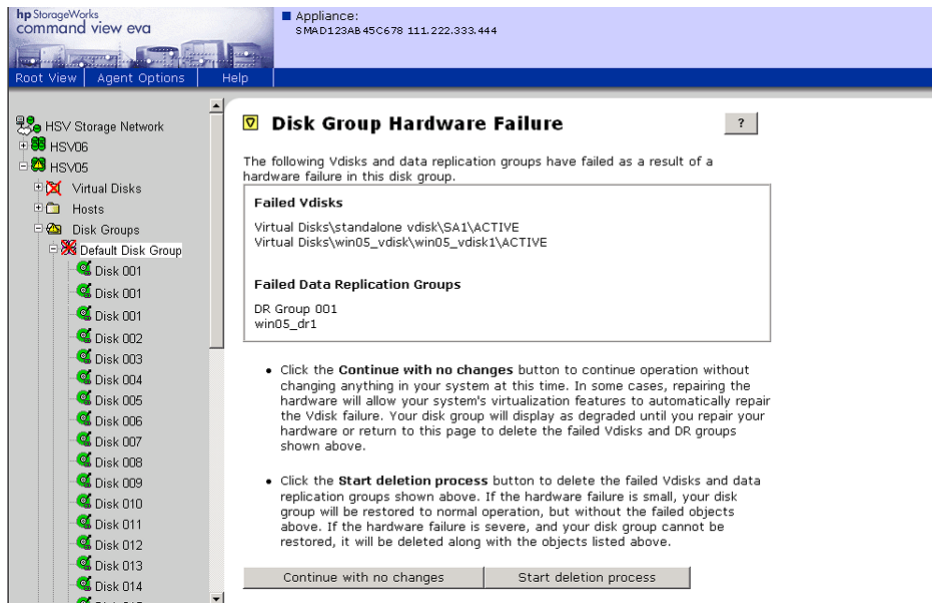


Figure 8 Disk Group Hardware Failure window

4. Click the Start deletion process tab.

A message displays requesting confirmation.

5. Click OK.

A list of affected DR groups requiring deletion is displayed.

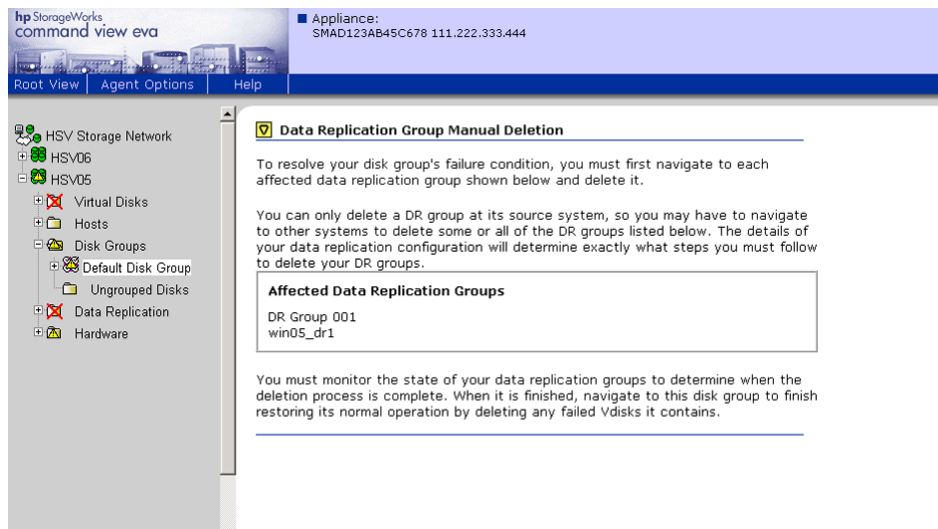


Figure 9 Data Replication Group Manual Deletion window

6. Select an affected DR group and click Delete.
A message is displayed to inform you that a DR group is being deleted.
7. Click OK.
The affected DR groups are deleted.
8. Select the failed virtual disks that were members of the affected DR groups.
A message is displayed while virtual disks are being deleted.

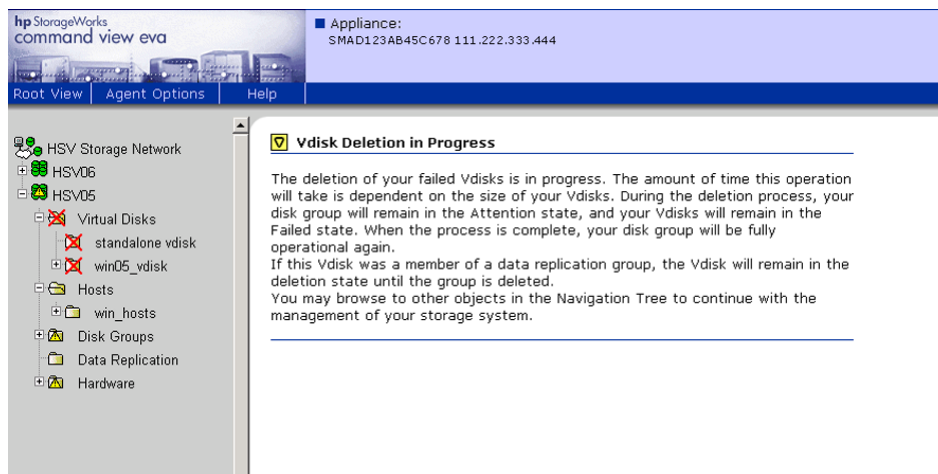


Figure 10 Vdisk Deletion in Progress window

When the deletion completes, an HP Command View EVA virtual disk Folder Properties screen is displayed, showing the virtual disk was deleted.

9. Navigate to the disk group and click Finish.
10. On the surviving array, delete the source DR group associated with the failed DR groups deleted in step 6.
11. (Optional) Repair your hard drives and re-create your disk group (see the HP Command View EVA documentation).

12. In the replication manager on the surviving array, re-create the DR groups and set up host presentation on the repaired array.
13. After normalization occurs between the source and destination arrays, fail over to the repaired array using the procedure described in [Planned failover](#).

Recovery when data replication was logging before failure

If data is logging when a source disk group hardware failure occurs, the data on the destination array is stale (not current). You have the following options:

- Recover using the stale data on the destination array (see [Disk group hardware failure on the source array](#)).
- Recover from a known, good point using a backup.
- If you want to perform a failover to quickly activate the destination array before repairing the inoperative disk group, use the procedure on [Disk group hardware failure on the source array](#).
- If you want to repair the inoperative disk group first, perform the repair, delete the inoperative DR groups and virtual disks on the failed system, re-create your virtual disks and DR groups, then restore your data from an external backup.

Disk group hardware failure on the destination array

This section describes how to recover from an inoperative disk group on your destination array. Your first indications that a disk group has become inoperative may be icons like those shown in [Table 6](#), except that your destination disk group status is Unknown.



NOTE:

When you delete DR groups to recover from a disk group hardware failure, you lose the redundancy of the other site or disaster tolerance of your data.

Perform this procedure using HP Command View EVA when a disk group hardware failure occurs on the destination array and the data on the source array is current.

1. In HP Command View EVA, begin troubleshooting the disk group problem.
2. Navigate to the failed disk group.

A list of failed virtual disks and DR groups is displayed.

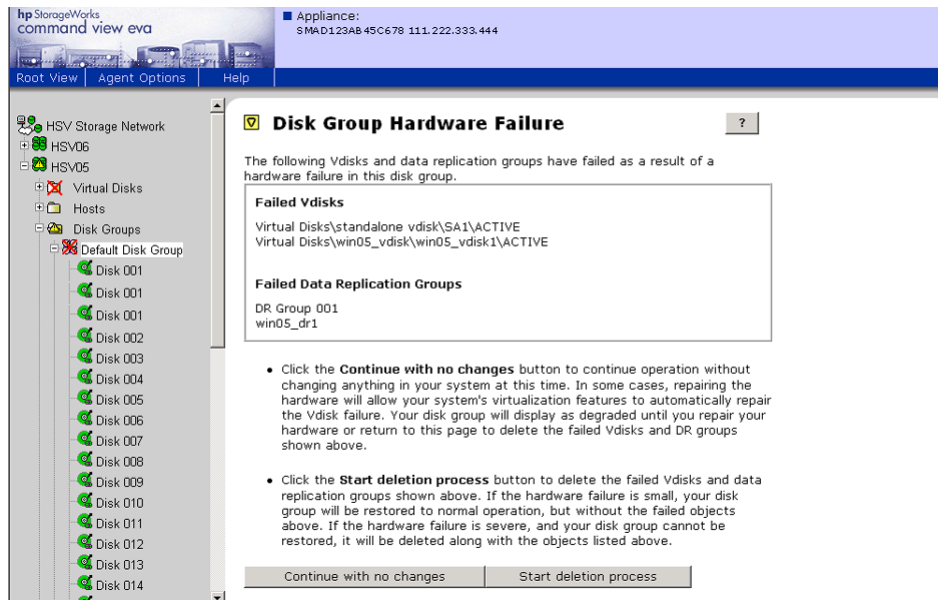


Figure 11 Disk Group Hardware Failure window

3. Click the Start deletion process tab.

A confirmation message is displayed.

4. Click OK.

A list of affected DR groups requiring deletion is displayed.

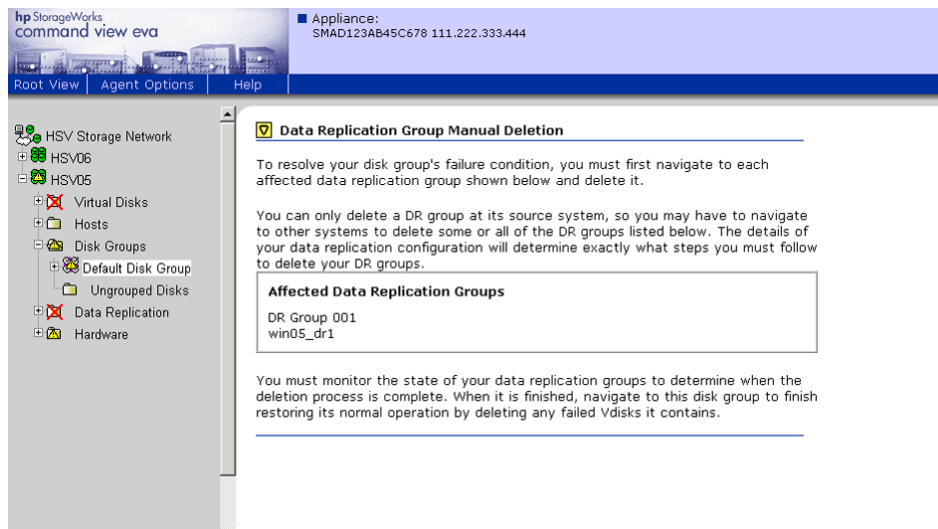


Figure 12 Data Replication Group Manual Deletion window

5. Select an affected DR group and click Delete.

A confirmation message is displayed.

6. Click OK.

The affected DR group is deleted.

7. Repeat [Step 5](#) and [Step 6](#) for each affected DR group.

8. Select failed virtual disks that were members of the affected DR groups.

A message is displayed while virtual disks are being deleted.

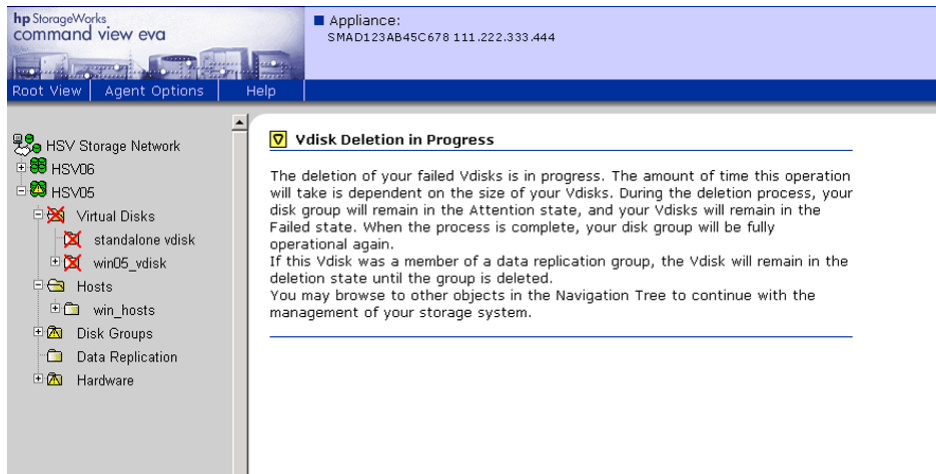


Figure 13 Vdisk Deletion in Progress window

When the deletion completes, an HP Command View EVA virtual disk Folder Properties screen is displayed showing the virtual disk was deleted.

9. Navigate to the disk group and click Finish to resolve the disk group hardware failure.
10. On the surviving array, delete the source DR group associated with the failed DR groups deleted in [Step 6](#).
11. (Optional) Repair your hard drives and re-create your disk group (see the HP Command View EVA documentation).
12. In the replication manager on the surviving array, re-create the DR groups and set up host presentation on the repaired array.

4 Operating System Specifics

Overview

This chapter provides operating system specific information related to stopping host I/O on the source array and the presentation of units to remote hosts to start host I/O.

Stopping host I/O

Follow the steps listed below for each operating system in your heterogeneous configuration. To stop all host I/O on the source array:

- HP OpenVMS—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and unmount the volumes associated with these virtual disks.
- HP Tru64 UNIX—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and unmount the volumes associated with these virtual disks.
- HP-UX—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, then unmount the file systems associated with the virtual disks.
- IBM AIX—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, then unmount the file systems associated with the virtual disks.
- Linux—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, then unmount the file systems associated with the virtual disks.
If you are running Logical Volume Manager (LVM) with or without clustering, see [Bootless DR group planned failover with Linux](#).
- Microsoft® Windows® NT–X86—If the operating system is up and running, shut it down.
- Microsoft Windows 2000/2003—If the operating system is up and running, shut it down.
- Novell NetWare—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, then dismount the volumes associated with these virtual disks.
- Sun Solaris—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and unmount the volumes associated with these virtual disks.

Allowing remote host discovery of devices

This section describes the commands you issue for each OS to allow remote hosts to discover devices.

HP OpenVMS

For HP OpenVMS, issue the following commands to allow remote host discovery of devices.

1. If the remote hosts are shut down, boot them now. Booting the hosts enables OpenVMS to recognize the drives.
2. If the remote hosts are not shut down, use the following command from a privileged account to enable OpenVMS to recognize the drives: `MCR SYSMAN IO AUTOCONFIGURE/LOG`

HP Tru64 Unix

For HP Tru64 Unix, issue the following commands to allow remote host discovery of devices.

1. If the remote hosts are shut down, boot them now. Booting the hosts enables Tru64 Unix to recognize the drives.
2. If the remote hosts are not shut down, use the following command to recognize the drives:
`hwmgr -scan scsi`

This may take a while for large configurations. If this is the case, scan only the SCSI buses that have new units added. Scan only one bus at a time. Use the following command:

```
hwmgr -scan scsi -bus x
```

(where x is the SCSI bus number)

HP-UX

For HP-UX, issue the following commands to allow remote host discovery of devices.

1. If the remote hosts are shut down, boot them now. Booting the hosts enables HP-UX to recognize the drives. If the system is already up, you will need to execute the following command to scan for the new devices:

```
# ioscan -fnCdisk
```

2. If the device special files are not present you would need to execute the following. This will add the special files.

```
# insf -e
```

A subsequent `ioscan` will return the devices with the special files.

- Once the system is up (and disks / special files are present), you need to create the folders for the new `vgxx` devices. Create the initial directory using `mkdir` command, then make the special file using the `mknod` command.

```
# mkdir <volumeGroupName>
```

```
# mknod <volumeGroupName> group c 64 minor#
```

Example:

```
# mkdir /dev/vg09
```

```
# mknod /dev/vg09/group c 64 0x090000
```

- Once the folders are created, the disks show up, and special files are present, you can then use the `vgimport` to create the virtual group.

```
# vgimport <volumeGroupName> <deviceSpecialFiles>
```

Example:

```
# vgimport /dev/vg09 /dev/dsk/c18t0d /dev/dsk/c18t1d0  
/dev/dsk/c25t0d0
```

- You can then attempt to display this virtual group using `# vdisplay -v /dev/vg09`. If this returns an error about "Volume group not activated" then you will need to activate it using the `vgchange` command.

```
# vgchange -a y <volumeGroupName>
```

Example:

```
# vgchange -a y /dev/vg09
```

- You may receive errors trying to mount the failed over volume (an error stating the volume is corrupt). You will need to run a file system check (this is fairly typical as the file system may not have been properly dismounted). You can repair the device using the `fsck` command.


```
# fsck <logicalVolumeName>
```

Once the devices are clean, the devices can be mounted.



NOTE:

VolumeGroupName is the name of the volume group you originally created at the local site. The *DeviceSpecialFiles* are from the ioscan in the form of /dev/dsk/c_t_d/. For consistency, configure the same *DeviceSpecialFiles* with the same volume groups, logical volumes, and file systems for the failed-over LUNs at the remote site with the same LUNs at the local site.

IBM AIX

For IBM AIX, issue the following commands to allow remote host discovery of devices.

1. If the remote hosts are shut down, boot them now. Booting the hosts enables AIX to recognize the drives.
2. If the remote hosts are not shut down, use the following commands to recognize the drives and verify that they are present:

```
cfgmgr -v
```

```
lsdev -Cc disk
```

Use the following commands to access file systems on the failed-over virtual disks:

```
importvg -y VolumeGroupName hdiskx
```

```
mount all
```



NOTE:

VolumeGroupName is the name of the volume group you originally created at the local site, and *x* is the number of the hdisk assigned to the failed-over virtual disk. If the *-y VolumeGroupName* parameter is omitted, AIX creates a default volume group name for you (for example, *vg00*) .

Linux

For Linux, perform the following tasks to allow remote host discovery of devices.

1. Reboot the servers at the remote site.
2. Remount the file system.

Microsoft Windows NT

For Microsoft Windows NT, perform the following tasks to allow remote host discovery of devices.

- Reboot the servers at the remote site and log on using an account that has administrative privileges. You should be able to see all of the units in My Computer.

Microsoft Windows 2000/2003

For Microsoft Windows 2000/2003, perform the following tasks to allow remote host discovery of devices.

1. On each host, log on using an account that has administrative privileges.
2. Open Computer Management and click **Disk Management**.
3. After Disk Management has initialized, select **Actions > Rescan Disks**. If the units fail to appear, click **F5** (Refresh). All of the failed-over units are displayed in the right pane.

Novell NetWare

For Novell NetWare, perform the following tasks to allow remote host discovery of devices.

- If the remote hosts are shut down, boot them now. If you are using traditional NetWare volumes, booting the hosts allows Novell NetWare to recognize the drives and automatically mount the volumes. If you are using NSS logical volumes, booting the hosts will recognize the NSS pools and activate them. However, you must manually mount each individual NSS volume by entering `MOUNT VolumeName` at the NetWare console.
- If the remote hosts are already up and running, or if they do not recognize the drives, issue the following command from the console before mounting the volumes:

```
SCAN FOR NEW DEVICES
```

Alternatively, you can use the `NWCONFIG` utility to issue this same command.

Next, mount the volumes with the following commands:

```
MOUNT ALL (for traditional NetWare volumes)
```

```
MOUNT VolumeName (for NSS logical volumes)
```

Sun Solaris

For Sun Solaris, perform the following tasks to allow remote host discovery of devices.

- Reboot the remote hosts using the `reboot -- -r` command, or use the following version-dependent commands to update the Secure Path Manager:
- Solaris 6, 7, and 8–


```
drvconfig -v
```

```
disks
```

```
/opt/CPQswsp/bin/spmgr display
```
- Solaris 9–
 - Present new units with LUN numbers sequentially following the old LUNs.
 - Run the following commands:


```
devfsadm -C
```

If you are using Secure Path:

```
/opt/CPQswsp/bin/spmgr display
```
- If Secure Path was not configured for these units, use the following version-dependent commands to add them to the Secure Path Manager.
 - Solaris 2.6, 7, and 8 –


```
/opt/CPQswsp/bin/spconfig
```

```
/opt/CPQswsp/bin/spmgr/display -u
```

```
/opt/CPQswsp/bin/spmgr add <WWLUNID>
```

```
devfsadm -C
```

```
/opt/CPQswsp/bin/spmgr display
```
 - Solaris 9 –

- Add the units with `spmgr add <WWLUNID>` or `spmgr add-r WWNN all`.
- Run `update_drv -f sd` to inform the system about attribute changes to the `sd` driver.
- Run `devfsadm -C` to create `/dev` entries for the new units.
- If you are using the `mpio`, and it has not been configured for these devices, issue the following commands to configure the new devices:

- Use the following `cfgadm` command as follows to list of the paths to the LUNS.

```
# cfgadm -al -o show_FCP_dev
```

Output:

```
c3 fc-private connected unconfigured unknown
```

```
c3::210000e08b0a5b65 unknown connected unconfigured unknown
```

```
c3::500060e802eb2b0b,0 disk connected unconfigured unknown
```

```
c4 fc-private connected unconfigured unknown
```

```
c4::210100e08b2a5b65 unknown connected unconfigured unknown
```

```
c4::500060e802eb2b14,0 disk connected unconfigured unknown
```

- The `cfgadm` command is used to configure the controller instances for `mpxio` control. Run this command for each EVA port on the storage array.

Example:

```
# cfgadm -c configure c3::500060e802eb2b0b
```

```
# cfgadm -c configure c4::500060e802eb2b14
```



NOTE:

The controller instance (c#) may differ between systems.

- If you are using Solaris 9, run this command to update the `sd` driver:
`#update_drv -f sd`
- Run the `devfsadm` command to build the appropriate device files:
`#devfsadm -C`
- If you are using Solaris 2.6, 7, or 8, reboot the host to finish configuring the new devices with the following command:
`# reboot - -r`
- You can now view the drives using the `format` command. See the current version of the multipath driver documentation, located at <http://www.sun.com/storage/san> for additional assistance.
- If you are using a volume manager application, import your volume groups. See your volume manager documentation.
- You may need `fsck` to repair any devices that were improperly dismounted.

5 Troubleshooting

Overview

This chapter provides troubleshooting guidance for arrays and links between multiple sites.

LUN inaccessible to host

A "stalled LUN" event (4206001b) in HP Command View EVA indicates that a LUN has been inaccessible to the host for at least four minutes, causing the LUN to be in a quiesced state. Take the following actions to troubleshoot this situation and prevent possible data loss:

1. Verify that the host still cannot access the LUN.
2. Try to resynchronize the controller from the HP Command View EVA field service page.
3. If the situation still exists, un-present and re-present the LUN to the host.
4. If the situation still exists, restart the controller and its partner controller, if necessary.

DR groups in unknown state

If your DR groups are in an unknown state, check to see if you have recently imported the replication manager database from an active management server to the management server where the DR groups are in an unknown state. If so, the problem is probably that the passwords do not match on the management servers.

Tunnel thrash

Tunnel thrash is the frequent closing and re-opening of a tunnel while holding host I/O in the transition. This occurs when peer controllers can see each other, but cannot sustain replication data with any path, even when throttled to the minimum. Some possible causes of tunnel thrash are:

- High volumes of packet loss
- Incorrectly configured routers
- Re-routed IP circuits
- Oversubscribed circuits

Although tunnel thrash is rare, a critical event (c23670c) is generated and displayed in HP Command View EVA for each DR group that shares the affected tunnel. You must intervene to prevent possible data loss.

Take the following actions to resolve this situation:

- Check all routers and look for high volumes of packet loss.
- Ensure that all router are configured correctly.
- Contact your service provider to check if the circuit has been alternate routed.
- Check to see if thrashing occurs during peak times and not during low volume times. If so, the circuit may be over subscribed and you may need to increase bandwidth.

**NOTE:**

An informational event (c22000c) is generated for an open tunnel. No action is required.

Remote server cannot detect a destination LUN

If you have a remote server that cannot detect a destination LUN, it could be that the DR group access mode is set to "disabled." A remote server can detect a LUN with a "read-only" access mode, but cannot detect it if the mode is set to "disabled." The replication manager allows you to change the DR group's access mode from "disabled" to "read-only," thereby allowing the remote server to detect the destination LUN. See the online help for information on editing a DR group's properties.

Long delays or time-outs on HP-UX

If an HP-UX host has multiple disk devices with failed or no longer presented LUNS behind them, it can take an increasingly long time to gather host information as the number of disk devices increases. If an HP-UX host exhibits time-outs on host discovery or failed jobs while waiting for host operations to complete, take the following actions:

- Check the disk devices showing long time-outs. Secure Path can display the status of the disk devices it is managing. For disk devices not managed by Secure Path, check for I/O time-outs by running an OS tool such as `diskinfo` on each disk device.

Remove any disk devices that show long time-outs, if they are no longer needed.

- If the disk devices are intentionally in this state, improve performance by modifying the I/O time-out setting for those disks with the `pvchange -t` command. HP-UX has a default I/O time-out of 30 seconds for SCSI disks. The `pvchange -t` command allows you to reduce the amount of time before a time-out on a given disk occurs. Reducing the time-out decreases the amount of time a host discovery takes.

6 Best practices

Overview

This chapter describes replication best practice procedures. It covers creating and using snapclones, bootless DR group planned failovers with Linux or SuSE Linux, throttling merge I/O after logging, and other miscellaneous best practices.

Creating a destination snapclone before making a full copy



NOTE:

A Business Copy EVA license is required for the following procedure.

When logging occurs on a source array, a temporary disparity occurs between data being processed at the local site and the data that exists at the remote location. A merge or full copy corrects this disparity later when the reason for the interruption is remedied. A merge sends data from the log disk in write order so it remains crash consistent. However, this is not the case with a full copy.

When a log fills to the point where it is marked for a full copy, there is a risk to the destination copy of the data once the process begins. This risk is due to the nature of a full copy, which copies data in 1-MB chunks starting at the beginning of the source virtual disk. This full-copy process does not maintain write ordering, and if it is not completed due to a second error, such as a loss of the source array, it leaves the destination array in an indeterminate state. Therefore, to prevent loss of data, best practice suggests creating a snapclone of destination virtual disks containing critical or important data prior to starting a full copy. If a major failure occurs at the local site during a full copy, the snapclone provides a clean copy of data as it existed before full copy writes were started to the destination array. However, any new writes that occurred on the source between the time the snapclone was created and the major failure would result in the loss of the new writes.

The following procedure describes the steps to take in a situation where you lose the connection between a source and destination array, and want to protect against a second failure when performing a full copy. Best practice suggests the creation of a destination snapclone whenever the link outage is expected to last more than several minutes.



NOTE:

You cannot use this procedure if a full copy has been started.

1. Using the replication manager, navigate to each affected DR group and suspend replication.
2. When able, use the managing server to make a snapclone of the destination virtual disks, using the procedures described in HP StorageWorks Business Copy documentation.
3. Using the replication manager, navigate to each affected DR group and resume replication. This will only enable replication if the links are still down.

Data movement using a snapclone



NOTE:

A Business Copy EVA license is required for the following procedure.

Use this procedure to move a copy of your data residing on a virtual disk to a remote location by the use of a snapclone. A snapclone is an exact copy of your virtual disk at the particular point-in-time it was created. The virtual disk being copied to the remote site becomes part of a DR group that can then be used as a new source virtual disk. You can use this procedure with data movement services such as:

- Data distribution—Pushing copies of data to other geographic locations to make it locally accessible.
- Data migration—Moving data to a new location or to one with a larger storage capacity.

To move data using a snapclone:


1. Make a snapclone of the virtual disk containing the data to be moved. See the online help for procedures on creating snapclones.

After the snapclone is created, the link from the snapclone to its original virtual disk dissolves, and the snapclone becomes a separate virtual disk.

2. Create a DR group with the new snapclone-created virtual disk linked to the remote array where you want the data to reside. The creation of the DR group replicates the virtual disk to your desired destination.
3. Delete the source-cloned virtual disk. You have the option of keeping the remote virtual disk or deleting the remote virtual disk.
4. Choose to keep remote virtual disk.

The data now resides as a new virtual disk on the remote array. It can be used as a source for another DR group, subject to the restriction that an array can be involved in a replicating relationship with only one other array.

Figure 14 provides a high-level summary of the following steps that perform data movement using a snapclone.



CXO8068b

Figure 14 Creating a DR group from a snapclone

Callouts:

1. HSV05 array
2. HSV06 array
3. HSV18 array
4. DR group
5. Virtual disk 1
6. Virtual disk 2
7. Replication
8. Virtual disk snapclone

Manually specifying disk group membership for a log-EVA 3000/5000 only

During the creation of a DR group, a 139-MB log is automatically created for the source and destination arrays. Placement of the log into a disk group on each system is based on the amount of free space in the disk group, the number of other logs in each disk group, the potential size of each existing log, and the potential size of the new log. It is possible that the disk group automatically selected for the

log may not be the same disk group where the DR group resides. For example, disk groups containing near-online FATA drives are automatically selected for the location of the log. If near-online drives are not present, disk groups with the largest amount of average free log space are chosen. The location of the log is rarely an issue, but you can control the placement of a log into a specific disk group to separate the log from the data.

By default, a log is created in a near-online disk group. You cannot force a log into an online disk group if a near-online disk group exists. This restriction applies to both the source and destination arrays.

If possible, creating DR groups and specifying log disk location should be done before other non-DR virtual disks are created. Doing this ensures sufficient available space in the disk groups for a log.

If there is enough available space in the destination array to create the destination virtual disks and log disk, then use the following procedure to manually specify the disk group where a log will be created:

1. Create Vraid0 virtual disks on the destination array in all disk groups except the one that will be used for the DR group destination virtual disk and log, so that all available space is filled. These virtual disks are temporary and are not used for data storage.
2. Select each disk group to check the available space on the destination array. Each disk group should report zero space available except the disk group where the DR group destination virtual disk and log will be created.
3. Wait until the virtual disks have finished allocating space and are in a normal state.
4. Create Vraid0 virtual disks on the source array in all disk groups except the one that will be used for the DR group source virtual disk and log, so that all available space is filled. These virtual disks are temporary and will not be used for data storage.
5. Select each disk group to check the available space on the source array. Each disk group should report zero space available except the disk group where the DR group source virtual disk and log will be created.
6. Wait until the virtual disks have finished allocating space and are in a normal state.
7. Create the DR group. Specify the disk group on the destination array where you want the destination virtual disk and log created. Do not allow the option to automatically select.
8. Delete the Vraid0 virtual disks that were created in step 1, then delete those created in step 4.

The log disk is now located in the disk group you selected. The disk group membership of the log is not visible to HP Command View EVA or to the replication manager interface.



NOTE:

With log placement, this procedure is not needed for EVA 4000/6000/8000 systems.

Three-site cascaded data replication using snapclones



NOTE:

A Business Copy license is required for the following procedure.

This procedure allows you to move copies of your data to a second remote location using HP Command View EVA and snapclones. The remote location can be an array without a replicating relationship to the array where the data was created. Exact copies of the virtual disks containing the data are created by using snapclones, and these are placed into a DR group for movement to the remote system.

For example, in [Figure 15](#) a production environment contains a DR group that replicates between arrays at Site 1 and Site 2. The DR group contains two virtual disks (05–06vdisk1 and 05–06vdisk2) that are to be archived on another array (Site 3). A snapclone of each virtual disk is created on the Site 2 array. After presentation to a host (set up only for presentation purposes, but required for the creation of a DR group), these members are added to a DR group called DR snapclone1. This DR group now resides on a source array that replicates to the desired destination array (Site 3). At the remote location, you can remove the virtual disk members from the DR group, renamed, and archived.



NOTE:

For this procedure, Site 1 is called the source array, Site 2 (the destination for the DR group from Site 1) is called the intermediate array, and Site 3 is referred to as the remote array.




Figure 15 Data movement using snapclones example

Callouts:

1. HSV05 array
2. HSV06 array
3. HSV18 array
4. DR group
5. Virtual disk 1
6. Virtual disk 2
7. Replication
8. Virtual disk snapclone

You can perform cascaded replication using any of the following methods:

- Using the job template in the replication manager
- Manually in the replication manager
- Manually with the command line user interface (CLUI)

The following procedure describes the steps you must perform, regardless of the method you use. See the replication manager online help and the *HP StorageWorks Replication Solutions Manager Command Line User Interface reference guide* for additional instructions.



NOTE:

An RSM job template is available for cascaded replication.

Before you begin

Before you begin the procedure, ensure that you have done the following tasks:

1. Set up DR groups at the source (site 1 in [Figure 15](#)) and the destination (site 2 in [Figure 15](#)).
2. Set up the host on the intermediate site (site 2 in [Figure 15](#)) using HP Command View EVA:
 - In the Add a Host window, enter a host name in the Host name box and a WWN in the Port WW Name box. Click the Add host tab.

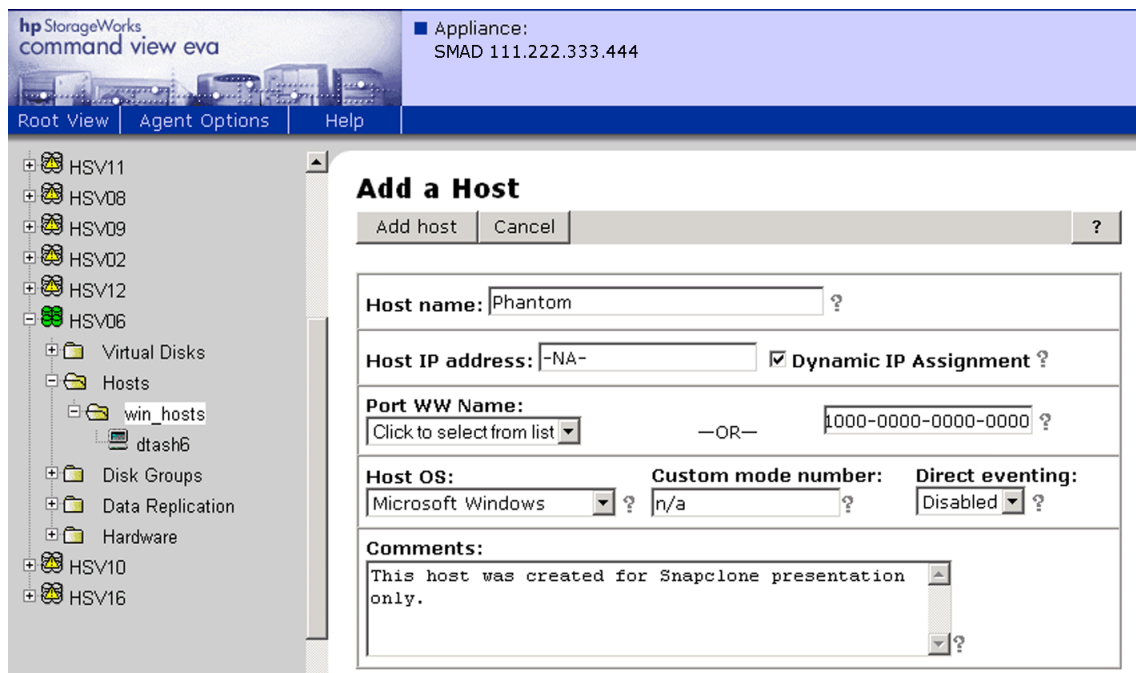


Figure 16 Add a Host window

An Operation completed page is displayed.

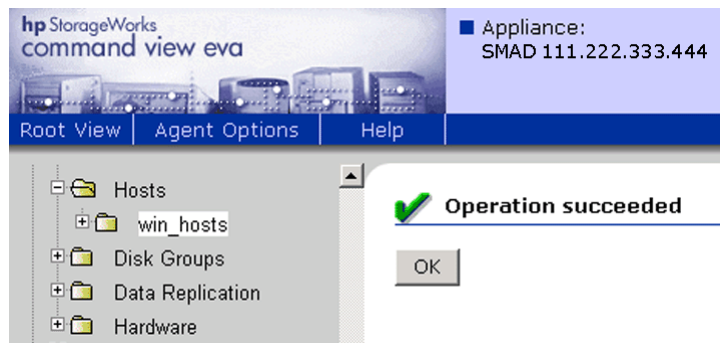


Figure 17 Operation succeeded page

Procedure

1. Enable failsafe mode for any DR group containing more than one replication pair.
2. Set synchronous write mode for any DR groups in this procedure.
3. If normalization is occurring to members of the DR group to be moved, wait for the members to normalize.
4. If an application requires that I/O be suspended before creation of a snapclone, suspend I/O at this time.
5. Create a snapclone of each virtual disk on the intermediate array (Site 2 in [Figure 15](#)).
6. Set the DR group back to asynchronous, if applicable.
7. Set failsafe mode back to Disabled for the DR group, if applicable.
8. If the application was suspended in step 4, restart the host application.
9. Present the snapclone(s) to the host on the intermediate site (site 2 in [Figure 15](#)).
10. Place the snapcloned virtual disks into a new DR group.
11. Wait for normalization to complete.
12. Unpresent the host from the snapcloned virtual disks in the DR group on the intermediate array (Site 2 in [Figure 15](#)).
13. Remove the remaining snapcloned virtual disks from the DR group by deleting the DR group. Leave the remote virtual disks intact by not discarding them during the deletion.
14. Delete the remaining snapcloned virtual disks from the intermediate array (Site 2 in [Figure 15](#)).

Post procedure

1. In HP Command View EVA, on the remote array (Site 3 in [Figure 15](#)), change the write protection of the mirrored snapcloned virtual disks to that of no write protection.
2. Rename the virtual disk to a useful name in the HP Command View EVA interface.

The virtual disks are now available on the remote array for any purpose.

Bootless DR group planned failover with Linux using LVM in standalone mode or with SuSE SLES 8 running LifeKeeper 4.4.3

The following procedures describe how to perform a bootless DR group failover when running the Logical Volume Manager (LVM) with Linux. Separate procedures for running in standalone host mode or with clusters (LifeKeeper) are listed. Perform the procedures for the source host, followed by the procedures for the destination host.



NOTE:

This procedure is not supported for unplanned failovers. The term "bootless" means that after the LUNs are first presented to a destination host, which requires an initial reboot, no further reboot of that host should be required.

Source host procedure

Perform one of the following steps on the source host, depending on whether or not you are running LifeKeeper 4.4.3.

1. If you are running LifeKeeper 4.4.3, proceed to step 2. If you are not running LifeKeeper, perform the following steps:
 - From your source host, stop I/O to your LUNs. Allow enough time for the I/O to complete before proceeding to the next step.
 - Unmount the volumes contained in the DR group.
Example: `umount /mounts/lvol1`
 - Change the status of the LUNs to inactive with the following command:
`vgchange VolumeGroupName -a n`
Example: `vgchange vg01 -a n`
 - Make the group unknown to the system with the `vgexport` command.
Example: `vgexport vg01`
 - Perform a failover of the DR group using the HP Continuous Access EVA user interface.
 - Depending on the number of LUNs, do one of the following to prevent Secure Path from detecting a failed disk:
 - For individual LUNs, run `spmgr quiesce -p path` (for each path visible to the LUNs).
 - For all LUNs at once, run `spmgr set -p off`. This method will turn off path verification for all LUNs still visible to the system.
2. If you are running LifeKeeper 4.4.3 clusters:
 - Bring your resources "out of service" with the LifeKeeper GUI.
 - From the system console:
 - Enter the `mount` command to verify the volume is unmounted.

- Enter the `vgscan` command to verify that the volume group is exported.
- Perform a failover of the DR group with the HP Continuous Access EVA user interface.
- Depending on the number of LUNs, do one of the following to prevent Secure Path from detecting a failed disk:
 - For individual LUNs, run `spmgr quiesce -p path` (for each path visible to the LUNs).
 - For all LUNs at once, run `spmgr set -p off`. This method will turn off path verification for all LUNs still visible to the system.

Destination host procedure

If this is the first time that LUNs are being presented to the destination host, reboot the host to pick up the new LUNs. If a reboot is not required (LUNs have been previously presented), and the paths are quiesced, issue the `spmgr restart all` command to unquiesce the paths.

Perform one of the following steps on the destination host, depending on whether or not you are running LifeKeeper 4.4.3.

1. If you are running LifeKeeper 4.4.3, proceed to step 2. If you are not running LifeKeeper, perform the following steps:
 - Issue the following command to make the volume known to the system:


```
vgimport VolumeGroupName PhysicalVolumePath
```

Example: `vgimport vg01 /dev/sda1`
 - Mount the file systems.

Example: `mount -t reiserfs /dev/vg01/lvol1 /mounts/lvol1`
 - Start host I/O.
 - If the verification path is turned off, issue the following command:


```
spmgr set -p on
```
2. If you are running LifeKeeper 4.4.3 clusters:
 - If this is the first time LUNs are being presented to the destination host, you must build the resource hierarchies for each new LUN presented (see the documentation on the LifeKeeper CD).
 - Bring your resources "out of service" with the LifeKeeper GUI.
 - Start host I/O.
 - If the verification path is turned off, issue the following command:


```
spmgr set -p on
```

Red Hat and SuSE Linux Lifekeeper clusters

Lifekeeper clusters must be zoned so that clustered hosts can see only one controller port per fabric. The operating system host mode of the controller must also be set to custom.

Throttling of merge I/O after logging

When I/O has been halted, DR groups not in failsafe-enabled mode automatically resume replication when links to the remote arrays are restored. If there are dozens of DR groups with large logs, they compete for bandwidth as they try to synchronize simultaneously.

By suspending the merging or copying of non-critical DR groups, the controllers merge only the most critical data first, allowing this data to be synchronized and become accessible before the less important data. As the more important groups finish merging, resume the I/O of the groups that were suspended. This concentration or channeling of I/O to specific groups by the use of suspend and resume commands is called throttling I/O.

Backing up replication jobs and configurations

HP recommends that you perform regular backups of jobs and configurations using the import and export features in your replication products. This ensures that job and configuration data can be easily restored during planned or unplanned maintenance of the replication server. See

Optimizing discovery refresh intervals

Use configuration settings in your replication products to optimize discovery refresh intervals. Replication products require up-to-date information on SAN resources. Find a balance between a discovery refresh interval that is too short (slowing overall performance with frequent discovery) and an interval that is too long (producing job failures due to out-of-date SAN information).

Optimizing browser-based GUI performance

Keep simultaneous browser sessions to the same replication manager to a minimum. A large number of sessions decreases responsiveness of the replication manager.

Coordinating enabled-host downtime

Ensure that planned downtime for enabled hosts is coordinated with replication jobs. A job will fail if any of a referenced host is not available when the job is run.

Minimizing simultaneous jobs

If you are using jobs, minimize simultaneous job execution, even if the jobs involve different arrays. For example, running too many replication manager jobs at the same time can reduce the overall responsiveness of the replication manager and of other applications on the replication management server.

Avoiding configuration changes while jobs are running

Avoid changing storage and host configurations while jobs are running. For example, don't change an array configuration with one interface (for example, HP Command View EVA) while replication jobs are running in another interface. Changing resources can lead to job failures and require manual intervention to restore resources to operational readiness.

Optimizing the number of active enabled hosts

Keep the number of active enabled hosts to the minimum needed to perform required operations. Consider stopping host agent processes on hosts when they are not needed for jobs. If operational changes result in a host no longer being used in jobs, consider removing the host agent. Reducing the number of host agents that communicate with the local replication server will result in better overall performance of the replication management server.

Coordinating enabled host shutdowns

Ensure that planned shutdowns are coordinated. Stopping an enabled host causes running jobs to fail if the job involves that host.

Coordinating replication server shutdowns

Ensure that planned shutdowns are coordinated. Stopping a replication server:

- Stops any local replication applications on the server.
- Causes running jobs to fail.
- Prevents scheduled jobs from starting.

Avoiding network identification changes

If possible, avoid changing the network identification (computer network name or IP address) of a replication management server or the computers on which replication host agents are running.

- Server identification change—If identification of a replication server is changed, use documented procedures to update all associated replication host agents so they can communicate with the replication management server. If the replication host agents are not updated to reflect the new replication server identification, jobs that involve the enabled host will fail.
- Host agent identification change—If identification of a replication host agent is changed, update impacted replication jobs so the jobs provide the correct references to enabled hosts. If the impacted jobs are not updated to reflect the new host agent identification, the jobs will fail.

Maintaining network connections

Ensure that network connections between a replication management server and enabled hosts are maintained, especially while jobs are running. Jobs that interact with enabled hosts fail if the network connection is not maintained throughout the job.

Using log files for troubleshooting jobs

The replication management server and host agents generate log files for job events. These detailed event log files can be helpful to HP support personnel when troubleshooting replication jobs.

Making CD-ROMs of replication product Web download files

HP recommends that you make installation and archive CD-ROMs of local replication files that you download from the HP Web site. Making copies of Web download files ensures that you can quickly

restore a replication management server and host agents to a given version without repeating download procedures.

Managing replication events

Develop operation guidelines and best practices to address the following situations and concerns.

Minimizing simultaneous replication events on an array

Minimize the number of replication requests to the same array at the same time. Consider limiting access to the various management and command line interfaces. Too many simultaneous replication events can reduce array performance.

Avoiding simultaneous replication events for the same virtual disk

Avoid making multiple replication requests to the same virtual disk at the same time. Multiple replication events to the same virtual disk not only slow performance, but in the case of automated jobs, can lead to job failures. For example, if the maximum number of snapshots per virtual disk is exceeded when the job is running, the job will fail.

Job scheduling

When using an external scheduler (or the internal scheduler in RSM) to schedule jobs, consider the timing of each job relative to other jobs that involve the array and host resources. Tune and load-balance demands to maximize performance.

Complying with EVA snapshot rules

The following general EVA snapshot rules apply:

- The array must have a local replication license.
- Each snapshot is created in the same disk family as the source virtual disk.
- All snapshots of a given virtual disk must have the same allocation policy.
- No more than seven snapshots of a given virtual disk can exist at one time.
- When managing array resources, snapshots are counted as a virtual disks.

Snapshots cannot be made when a storage volume:

- Is itself a snapshot.
- Is in the process of unsharing or being deleted.

When a local replication interface indicates that a EVA storage volume (virtual disk) does not support snapshot replication, or if a snapclone script action fails in a job, it is probable that some of these rules have been violated.

Complying with EVA snapclone rules

The following general EVA snapclone rules apply:

- The array must have a local replication license.
- Each snapclone is created in the same disk group as its source, but in a different disk family.

Snapclones cannot be made when a storage volume:

- Is itself a snapshot or has a snapshot.
- Is in the process of unsharing or being deleted.

When a local replication interface indicates that an EVA storage volume (virtual disk) does not support snapclone replication, or if a snapclone script action fails in a job, it is probable that some of these rules have been violated.

Caching in Microsoft Windows

Small files in Microsoft Windows can be held in cache, disrupting replication to the remote controller. Flush all cache files, if possible, before performing a failover. One source of information for flushing data caches on CPU and kernel architecture can be obtained from:

<http://msdn.microsoft.com/library/en-us/wcedsn40/html/cgconimplementingcacheflushroutines.asp>.

Another option is to use the HP StorageWorks Business Copy EVA application to flush the cache. For more information, go to:

<http://h18006.www1.hp.com/products/storage/software/bizcopyeva/index.html>.



NOTE:

Rebooting of the source host(s) is the only qualified procedure at this time.

Glossary

This glossary defines terms used in this guide or related to this product and is not a comprehensive glossary of computer terms.

active management server	See management server.
array	See virtual array and storage system.
asynchronous	A descriptive term for computing models that eliminate timing dependencies between sequential processes. In asynchronous replication, the array controller acknowledges that data has been written at the source before the data is copied at the destination. Asynchronous replication is an optional DR group property. See also synchronous.
bidirectional	A descriptive term for an array that contains both source and destination virtual disks. This configuration allows multidirectional I/O flow among several arrays
B-series switches	Fibre Channel core and SAN switches made by Brocade and sold by HP.
C-series switches	Fibre Channel switches made by Cisco and sold by HP.
client	An intelligent device that requests services from other intelligent devices. In the context of HP StorageWorks Replication Manager, a client is a computer that is used to access the replication manager remotely using a supported browser.
copy set	A source–destination pair of vdisks.
default disk group	The disk group that is created when an array is initialized. The minimum number of disks the group can contain is eight. The maximum is the number of installed disks.
destination	The targeted recipient (for example, a DR group, array, virtual disk) of replicated data. See also source.
disk group	A named group of disks selected from all the available disks in an array. One or more virtual disks can be created from a disk group.
DR group	Data replication group. A named group of virtual disks selected from one or more disk groups so that they replicate to the same destination, fail over together, and preserve write order within the group.
dual fabric	Two independent fabrics providing multipath connections between Fibre Channel end devices.
EVA	Enterprise Virtual Array, an HP StorageWorks product that consists of one or more virtual arrays. See also virtual arrays.
event	A system-generated status message, resulting from a: <ul style="list-style-type: none">• User-initiated action (for example, "suspend DR group")• Replication or system transaction (for example, "retrieved data for storage system")• Job operation (for example, "job complete")
fabric	A network of Fibre Channel switches or hubs and other devices.

failover	An operation that reverses replication direction so that the destination becomes the source and the source becomes the destination. Failovers can be planned or unplanned and can occur between DR groups, managed sets, fabrics or paths, and array controllers.
failsafe	A descriptive term for devices that automatically assume a safe condition after a malfunction. Failsafe DR groups stop accepting host input and stop logging write history if a member of the group becomes unreachable.
general purpose server	A server that runs customer applications such as file and print services. HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager can be used on a general purpose server in limited configurations.
HP Continuous Access EVA	HP Continuous Access EVA is a storage-based HP StorageWorks product consisting of two or more storage systems performing disk-to-disk replication, along with the management user interfaces that facilitates configuring, monitoring, and maintaining the replicating capabilities of the storage systems.
Home	The DR group that is the preferred source in a replication relationship. By default, Home is the original source, but it can be set to the destination DR group.
host	A computer that runs user applications and uses the information stored on an array.
host volume	Storage capacity that is defined and mountable by a host operating system. In HP Replication Solutions Manager, host volumes are disks or volumes that are reported by an enabled host.
initialization	A configuration step that binds the controllers together and establishes preliminary data structures on the array. Initialization also sets up the first disk group, called the default disk group, and makes the array ready for use.
LUN	Logical unit number. Logical units are the components within SCSI targets that execute I/O commands. Virtual disks that are presented to hosts correspond to logical units and are identified by LUN IDs. See also present.
M-series switches	Fibre Channel Director and Edge switches made by McDATA and sold by HP.
managed set	Selected resources grouped together for convenient management. For example, you can create a managed set to manage all DR groups whose sources reside in the same rack.
management server	A server where HP StorageWorks Enterprise Virtual Array (EVA) management software is installed, including HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager, if used. A dedicated management server runs EVA management software exclusively. Other management servers are general purpose servers, HP ProLiant Storage Server (NAS) models, and the HP OpenView Storage Management Appliance. When there are multiple management servers in a SAN, one is active and all others are standby. The active management server actively manages the array, while the standby management server takes control of the array if there is a failure on the active management server. There is only one active management server at a time for any given management zone in a SAN.
merge	The act of transferring log contents to the destination virtual disk to synchronize the source and destination.
mount point	The file system path and directory where a host volume is accessed.
normalization	The initial full copy that occurs between source and destination virtual disks.
(to) present	The array controller act of making a virtual disk accessible to a host computer.
remote copy	A replica virtual disk on the destination array.

resource	An object in the Replication Solutions Manager navigation pane; namely, DR groups, enabled hosts, host volumes, managed sets, storage systems, and virtual disks. Replication is performed using these resources.
SAN	Storage area network, a network of storage devices and the initiators that store and retrieve information on those devices, including the communication infrastructure.
snapclone	A copy that begins as a fully allocated snapshot and becomes an independent virtual disk. Applies only to the HP StorageWorks EVA.
snapshot	A nearly instantaneous copy of the contents of a virtual disk created without interruption of operations on the source virtual disk. Snapshots are typically used for short-term tasks such as backups.
source (Home)	A descriptive term for the virtual disk, DR group, or virtual array where an original I/O is stored before replication. See also destination.
source–destination pair (of vdisks)	A copy set.
standby management server	See management server.
Storage Management Appliance	HP OpenView Storage Management Appliance, an HP hardware–software product designed to run SAN management applications such as HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager.
storage system	Synonymous with virtual array. The HP StorageWorks Enterprise Virtual Array consists of one or more storage systems. See also virtual array.
synchronous	A descriptive term for computing models that perform tasks in chronological order without interruption. In synchronous replication, the source waits for data to be copied at the destination before acknowledging that it has been written at the source. See also asynchronous.
UUID	Unique Universal Identifier, a unique 128-bit identifier for each component of an array. UUIDs are internal system values that users cannot modify.
VCS	Virtual Controller Software. The software in the HP StorageWorks Enterprise Virtual Array controller. Controller software manages all aspects of array operation, including communication with HP StorageWorks Command View EVA.
virtual array	Synonymous with disk array and storage system, a group of disks in one or more disk enclosures combined with control software that presents disk storage capacity as one or more virtual disks. See also virtual disks.
Virtual Controller Software	See VCS.
virtual disk	Variable disk capacity that is defined and managed by the array controller and presentable to hosts as a disk.
Vraid	Techniques for configuring virtual disks to provide fault tolerance and increase performance. Vraid techniques are identified by level numbers. Level Redundancy Technique Vraid0 None Striping Vraid1 High Mirroring Vraid5 Medium Striping and parity
Vraid0	A virtualization technique that provides no data protection. Data chunks are distributed across the disk group from which the virtual disk is created. Reading

and writing to a Vraid0 virtual disk is very fast and uses available storage to the fullest, but provides no data protection (redundancy) unless there is parity.

Vraid1

A virtualization technique that provides the highest level of data protection. All data blocks are mirrored, or written twice, on separate disks. For read requests, the block can be read from either disk, which can increase performance. Mirroring requires the most storage space because twice the storage capacity must be allocated for a given amount of data.

Vraid5

A virtualization technique that uses parity striping to provide moderate data protection. For a striped virtual disk, data is broken into chunks and distributed across the disk group. If the striped virtual disk has parity, another chunk (a parity chunk) is calculated from the data chunks and written to the disks. If a data chunk becomes corrupted, the data can be reconstructed from the parity chunk and the remaining data chunks.

Index

A

- active site, 15
- alternate site, 17
- assumptions
 - EVA, 17
 - host operating systems, 18
 - hosts, 18
- audience, 9
- authorized reseller
 - HP, 12
- avoiding configuration changes while jobs are running, 72
- avoiding network identification changes, 73
- avoiding simultaneous replication events for the same virtual disk, 74

B

- backing up the configuration, 37
- best practices, 63, 74
 - avoiding configuration changes while jobs are running, 72
 - avoiding network identification changes, 73
 - avoiding simultaneous replication events for the same virtual disk, 74
 - backing up replication jobs and configurations, 72
 - bootless DR group failover (Linux), 70
 - coordinating enabled host shutdowns, 73
 - coordinating enabled-host downtime, 72
 - coordinating replication server shutdowns, 73
 - creating a destination snapclone before making a full copy, 63
 - maintaining network connections, 73
 - making CD-ROMs of replication product, 73
 - managing replication events, 74
 - minimizing simultaneous replication events on an array, 74
 - optimizing browser-based GUI performance, 72
 - optimizing discovery refresh intervals, 72
 - optimizing the number of active enabled hosts, 73
 - snapclone rules, 74
 - snapshot rules, 74
 - specifying disk group membership for a log, 65
 - three-site cascaded replication, 66
 - throttling of merge I/O after logging, 72
 - using a snapclone to move data, 64
 - using log files for troubleshooting, 73
 - Windows caching, 75
- best practices;
 - support procedures, 63
- bidirectional replication, 15, 26
- bootless DR group failover (Linux), 70
- Business Copy EVA

- license, 63, 64

C

- caching
 - Windows, 75
- capturing configuration information
 - manually, 38
 - using SSSU, 37
- cascaded replication, 66
- component repair vs. failover, 36
- concepts, 25
 - bidirectional replication, 26
 - DR group, 26
 - failover, 30
 - failsafe mode, 33
 - Home designation, 27
 - local replication, 25
 - remote replication, 25
 - replication direction, 26
 - snapclone, 25
 - snapshot, 25
 - zoning, 19
- configuration
 - backing up, 37
- configuration information
 - capturing manually, 38
 - capturing, using SSSU, 37
- configuration: text files
 - , 37
- conventions
 - document, 11
- coordinating enabled host shutdowns, 73
- coordinating enabled-host downtime, 72
- coordinating replication server shutdowns, 73
- creating
 - destination snapclone before full copy, 63

D

- data movement using a snapclone, 64
- data replication, 25
- data replication group, 26
- destination array, 15
- destination virtual disk, 25
- disaster planning, 35
- disk group
 - hardware failure, 49
 - hardware failure, definition, 49
 - hardware failure, on source array, 49
- disk group:hardware failure:on the destination array
 - , 52
- document conventions, 11

- DR group
 - bootless failover (Linux), 70
 - description, 26
 - failover, 36
 - log disks, 28
 - log size, 29
 - logging state, 28
 - presentation, 27
 - presenting DR group members to same HBA, 28
 - properties, 28
 - unknown state, 61

E

- Element Manager for HSG, 19
- Enterprise Virtual Array
 - description, 17
- EVA, 17
 - assumptions, 17
- event
 - logs, 33
 - scenarios, 39

F

- failover, 16
 - concept, 30
 - controller, 36
 - defined, 36
 - DR group, 36
 - fabric or path, 36
 - managed set, 36
 - planned procedure, 42
 - planned scenario, 36, 40
 - unplanned procedure, 44
 - unplanned scenario, 40
- failover vs. component repair, 36
- failover;
 - unplanned scenario, 36
- failsafe mode, 33
- fast synchronization, 28
- FATA drives, 65
- Fibre Channel adapter
 - in host, 22
- full copy, 28, 63

H

- HBA
 - presenting DR group members, 28
- Home designation
 - DR group, 27
- host operating systems, 18
 - assumptions, 18
- hosts
 - assumptions, 18
- HP Command View EVA
 - description, 19, 22
 - interface options, 20
- HP Continuous Access EVA

- failover, 35
- features, 16
- overview, 15
- prerequisites, 9
- related documentation, 9
- HP OpenVMS
 - privileges, 55
- HP Replication Solutions Manager
 - interface options, 21
- HP resources
 - authorized reseller, 12
 - storage website, 12
 - technical support, 12
- HP StorageWorks SMI-S
 - interface options, 22
- HP-UX time-outs and delays, 62

I

- inoperative disk group, 49
- interface options, 20
 - HP Command View EVA, 20
 - HP Replication Solutions Manager, 21
 - HP StorageWorks SMI-S, 22
 - SSSU, 22

J

- jobs
 - minimizing simultaneous, 72
 - scheduling, 74

L

- license keys, 17
- licensing, 17
- Lifekeeper clusters
 - Red Hat and SuSE Linux, 71
- local replication
 - concept, 25
- local site, 15
- log disks, 28
 - description, 28
 - fast synchronization, 28
 - full copy, 28
 - group membership, 65
- log disks:
 - marked for full copy, 63
- log size, 29
- logging, 50
- logging states, 28
- logs
 - event, 33
 - security, 33
 - trace, 33
 - transaction, 33
- long delays on HP-UX, 62
- loss of redundancy, 49
- LUN
 - inaccessible to host, troubleshooting, 61

read-only access for destination, 62

M

- maintaining network connections, 73
- making CD-ROMs of replication product, 73
- managed set, 29
 - failover, 36
- managing replication events, 74
- merging, 63
- minimizing simultaneous replication events on an array, 74
- mode
 - failsafe, 33

O

- optimizing browser-based GUI performance, 72
- optimizing discovery refresh intervals, 72
- optimizing the number of active enabled hosts, 73
- original state, 26
- overview
 - HP Continuous Access EVA, 15

P

- path failover, 36
- placement of log into disk group, 65
- planned failover, 36, 42
 - scenario, 40
- planning for disaster, 35
- prerequisites, 9
- presentation
 - DR group, 27
- primary site, 17

R

- read-only access
 - destination LUN, 62
- Red Hat Linux:
 - Lifekeeper clusters, 71
- related documentation, 9
- remote copy, 25
- remote replication
 - bidirectional replication, 26
 - concept, 25
- remote site, 15
- replication direction
 - concept, 26
- Replication Solutions Manager interface
 - managed sets, 29
- Resume command, 63
- resumption of operations procedure, 44
- resumption of operations scenario, 41
- return operations to Home array, 45
- return operations to Home array scenario, 41
- return operations to replaced new hardware, 41, 46

S

- security logs, 33
- setting read-only access for a destination LUN, 62
- single component failure, 36
- site failover
 - description, 36
- snapclone
 - concept, 25
 - creating before full copy, 63
 - data movement, 64
 - rules, 74
 - three-site cascaded replication, 66
- snapshot
 - concept, 25
 - rules, 74
- source array, 15
- source virtual disk, 25
- specifying disk group membership for a log, 65
- SSSU, 48
 - interface options, 22
- stale data, 50
- standby site, 15
- Storage System Scripting Utility; , 48
- support procedures, 63
- supported operating systems, 18
- SuSE Linux
 - Lifekeeper clusters, 71
- Suspend command, 63

T

- technical support
 - HP, getting help, 12
- three-site cascaded replication, 66
- throttling I/O, 72
- throttling of merge I/O after logging; , 72
- time-outs on HP-UX, 62
- trace logs, 33
- transaction logs, 33
- troubleshooting
 - DR groups in unknown state, 61
 - long delays or time-outs on HP-UX, 62
 - LUN detection, 62
 - tunnel thrash, 61
- troubleshooting: ; :
 - LUN inaccessible to host, 61
- tunnel thrash
 - troubleshooting, 61

U

- unknown state
 - DR groups, 61
- unplanned failover, 36, 40
 - procedure, 44
- using log files for troubleshooting, 73

V

- VCS, 18

Virtual Controller Software
description, [18](#)

[W](#)
websites

HP storage, [12](#)
technical support, [12](#)
Windows caching, [75](#)